arXiv:1409.0309v1 [cs.CR] 1 Sep 2014

Using Architecture to Reason About Information Security*

Stephen Chong
Harvard University

Ron van der Meyden
UNSW Australia

Abstract

We demonstrate, by a number of examples, that information-flow security properties can be proved from abstract architectural descriptions, that describe only the causal structure of a system and local properties of trusted components. We specify these architectural descriptions of systems by generalizing intransitive noninterference policies to admit the ability to filter information passed between communicating domains. A notion of refinement of such system architectures is developed that supports top-down development of architectural specifications and proofs by abstraction of information security properties. We also show that, in a concrete setting where the causal structure is enforced by access control, a static check of the access control setting plus local verification of the trusted components is sufficient to prove that a generalized intransitive noninterference policy is satisfied.

Contents

1 Introduction **2**

2 Architectures and semantics **4**
2.1 Example: $\mathcal{HL}$ architecture 5
2.2 Example: Hinke-Schaefer 5
2.3 Machine model 6
2.4 Semantics 7

3 Information security properties **8**
3.1 Group knowledge vs. Distributed knowledge 9
3.2 Knowledge and architectural refinement 11

*Version of September 2, 2014. An early version of this paper circulated in 2009 and has been cited in the literature. The present version includes a significant reworking of results on architectural refinement from that version. A version of the paper introducing the access control implementation results was presented at the Layered Assurance Workshop at ACSAC 2012. The present version adds to these earlier versions full proofs of all results and an extended set of examples.

3.3 Example: $\mathcal{HL}$ information security 14
3.4 Example: Hinke-Schaefer 14

4 Extended architecture **15**
4.1 Filter functions 15
4.2 Example: Starlight Interactive Link 21
4.2.1 Information security properties 23
4.3 Example: Downgrader 25
4.4 Example: Electronic election 29

5 Architectural refinement **32**
5.1 Architectural refinement for interpreted extended architectures . 33
5.1.1 Sufficient conditions for semantic refinement 34
5.1.2 Properties of architectural refinement 39
5.2 Specification Refinement 40
5.3 Example: Hinke-Schaefer 41
5.4 Example: Downgrader 41

6 Implementing Architectures using Access Control **46**
6.1 Starlight Architecture 51
6.2 Downgrader Architecture 54
6.3 Electronic Election 56

7 Enforcing Architectures on Concrete Platforms **58**

8 Related Work **60**

9 Conclusion **62**

1 Introduction

System architectures are high-level designs that describe the overall structure of a system in terms of its components and their interactions. Proposals for architectural modeling languages (e.g., AADL and Acme [Garlan et al., 2000]) vary with respect to their level of detail and contents, but at the most abstract level, architectures specify the *causal structure of a system*.

The MILS (Multiple Independent Levels of Security and Safety) initiative [Alves-Foss et al., 2006; Vanfleet et al., 2005; Boettcher et al., 2008] of the US Air Force proposes to use architecture as a key part of the assurance case for high-assurance systems. The details of the MILS vision are still under development but, as articulated by Boettcher et al. [2008], it encompasses a 2-level design process, consisting of a policy level and a resource sharing level.

At the policy level, the system is described by an architecture in the form of a graph, in which vertices correspond to components and the edges specify permitted communication between components. In this respect, the architecture is like an *intransitive noninterference* security policy [Haigh and Young, 1987;

Rushby, 1992; van der Meyden, 2007]. At the policy level, one might also specify which components are trusted, and the local policies that these components are trusted to enforce. According to the MILS vision, building a system according to the architecture, by composing components that satisfy their local policies, should result in the system satisfying *global* security and safety properties.

At the resource sharing level, MILS envisages the use of a range of infrastructural mechanisms to ensure that the architectural information-flow policy is enforced despite components sharing resources such as processors, file systems, and network links. These mechanisms might include physical isolation, separation kernels, periods processing, cryptography and separating network infrastructure. It is intended that this infrastructure will be developed to a high level of assurance, so that a systems assurance case can be obtained by the composition of the assurance cases for trusted components and systems infrastructure. It is hoped this will enable a COTS-like market for infrastructural mechanisms and trusted components.

The key contribution of this paper is to demonstrate, through several examples, that it is in fact possible, as envisaged in the MILS literature, to derive interesting information security properties compositionally from a high-level specification of trusted components and their architectural structure. We focus on compositional reasoning about information-flow security properties.

We present a framework that allows the specification of a system architecture with local constraints on some system components. To give a precise meaning to the architectural structure, we extend the semantics for intransitive noninterference developed by van der Meyden [2007]. An architectural interpretation of this semantics has previously been given [van der Meyden, 2012]. In order to express constraints on trusted components, we extend architectures by labeling edges between components with functions that further restrict the information permitted to flow along edges. One of the contributions of the paper is to give a formal semantics to the enriched architectures that include these new types of edges. We also develop a theory of refinement for these enriched architectures, which enables top-down, correctness-preserving development of architectural specifications. It also enables simple proofs of information security properties on complex architectures to be obtained using an abstraction of that architecture.

We demonstrate the use of the framework through examples motivated by systems with interesting security requirements. These include multi-level secure databases, the Starlight Interactive Link [Anderson et al., 1996], a trusted downgrader, and a simple electronic election system.

In each example, we identify an architectural structure and a mathematically precise set of local constraints on the trusted components. We then show that information-flow properties expressed in a logic of knowledge arise as a consequence of the interaction of the local constraints and the architectural structure. Our results show that for *any* system that is compliant with the architecture, if the trusted components satisfy their local constraints then the system satisfies the global information-flow properties.

The information security properties presented in the examples provide information-

theoretic and application-specific guarantees. Thus, there are no covert channels that can violate the information security properties, and the negation of each information security property would constitute an application-specific attack.

Only a few examples have been presented to date to formally justify the MILS approach to high-assurance secure systems development. One example is developed in Greve et al. [2003], but with respect to a more concrete model (based on a separation kernel formal security policy that deals with access control on memory segments) than the abstract, "noninterference" style semantics we consider. Our policy level model is more abstract, and allows greater flexibility for implementations. However, we also consider a more concrete model, systems with structured state subject to "reference monitor conditions" [Rushby, 1992]. We show that in this setting, to prove that a system complies with one of our extended architectures, it suffices to check a simple condition on the access control setting and to prove local properties of the trusted components.

By developing an abstract semantics for architectures and specifications of trusted components, and by developing additional examples, our work advances the case that global information-flow security properties can be derived from a high-level systems architecture and local constraints on trusted components within this architecture, in the style of reasoning envisaged by Boettcher et al. [2008].

The structure of the paper is as follows. In Section 2 we review architectures and their semantics. In Section 3, we introduce the epistemic logic we use to express information security properties. In Section 4 we extend architectures with *filter functions* that allow fine-grained specification of what information flows between components. The extended architectures enable the proof of additional information security properties. Section 5 extends previous work on architectural refinement to account for filter functions. The concrete model based on access control is developed in Section 6 and we consider possible platforms and techniques that might be used to show the access control model holds in Section 7. We discuss related work in Section 8. Section 9 concludes. We use examples throughout, to illustrate and motivate the definitions and results.

2 Architectures and semantics

Architectures give a policy level description of the structure of a system. We begin with a simple notion of architecture, following van der Meyden [2012]. A richer notion will be introduced later.

An architecture is a pair $\mathcal{A} = (D, \rightarrowtail)$, where D is a set of security domains, and the binary relation $\rightarrowtail \subseteq D \times D$ is an *information-flow policy*. The relation $\rightarrowtail$ is reflexive but not necessarily transitive. Intuitively, information is allowed to flow from domain u to domain v only if $u \rightarrowtail v$. The relation is reflexive as it is assumed that information flow within a domain cannot be prevented, so is always allowed.

In the literature on information flow policies, domains are generally understood to correspond to security levels. We use a more general interpretation,

in which domains can may also correspond to system components or agents in the system. In an implementation of an architecture, separate domains do not necessarily utilize separate resources. Hardware, code, and data may be shared between domains. Indeed, a key challenge is ensuring that the information-flow policy is respected despite the shared use of resources.

2.1 Example: $\mathcal{HL}$ architecture

The architecture $\mathcal{HL} = (\{H, L\}, \{(L, L),\ (H, H),\ (L, H)\})$ consists of two security domains H and L, and the information-flow policy indicates that information is allowed to flow from L to H, in addition to the reflexive information flows. We can depict $\mathcal{HL}$ graphically, indicating security domains with rectangles, and the information-flow policy with arrows. We omit arrows for reflexive information flows.

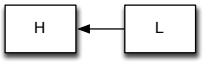

2.2 Example: Hinke-Schaefer

A variety of architectures have been proposed for multi-level secure database management systems (MLS/DBMS) [Thuraisingham, 2005]. In the Hinke-Schaefer architecture [Hinke and Schaefer, 1975], several (untrusted) single-level DBMSs are composed together in a trusted operating system. Each user interacts with a single-level DBMS. The operating system enforces access control between the single-level DBMSs, allowing more restrictive DBMSs to read the storage files of less restrictive DBMSs, but not vice versa.

The following diagram shows architecture $\mathcal{HS}$, which represents the Hinke-Schaefer architecture for two security levels at the MILS policy level.

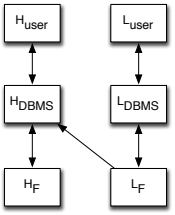

Domains H_{user} and L_{user} represent users of a high-security and low-security DBMS respectively; they interact with the single-level DBMSs H_{DBMS} and L_{DBMS} respectively. The single-level DBMSs store their data in database files denoted H_F and L_F. Note that information is allowed to flow to H_{DBMS} from both H_F and L_F, as the high-security DBMS is allowed to read the storage files of both the high-security and low-security DBMSs.

The Hinke-Schaefer architecture is also known as the "operating system providing mandatory access control" architecture [Thuraisingham, 2005], as the

operating system is trusted to enforce the information flows specified in the architecture. This amounts to a decision to implement the policy level architecture $\mathcal{HS}$ at the resource sharing level by means of a trusted separation kernel.

2.3 Machine model

To specify what it means for an implementation to satisfy an architecture, we must first define what an implementation is. We use the *state-observed* machine model [Rushby, 1992], which defines deterministic state-based machines. A machine has a set of actions A, and each action is associated with a security domain. Intuitively, if action a is associated with domain u, then a represents a decision, choice, or action taken by the system component represented by u. Actions deterministically alter the state of the machine, and we assume that the observations of each security domain are determined by the current machine state.

A machine is a tuple $M = \langle S, s_0, A, D, \mathtt{step}, \mathtt{obs}, \mathtt{dom}\rangle$ where S is a set of states, $s_0 \in S$ is the *initial state*, A is a set of *actions*, D is a set of domains, $\mathtt{step} : S \times A \to S$ is a deterministic transition function, $\mathtt{dom} : A \to D$ associates a domain with each action, and observation function $\mathtt{obs} : D \times S \to O$ describes for each state what observations can be made by each domain, for some set of observations O.

We assume that it is possible to execute any action in any state: the function $\mathtt{step}$ is total. Given sequence of actions $\alpha \in A^*$, we write $s \cdot \alpha$ for the state reached by performing each action in turn, starting in state s. We define $s \cdot \alpha$ inductively defined using the transition function $\mathtt{step}$, by

$$s \cdot \epsilon = s$$
$$s \cdot \alpha a = \mathtt{step}(s \cdot \alpha, a)$$

for $\alpha \in A^*$ and $a \in A$. (Here ϵ denotes the empty sequence.) For notational convenience, we write $\mathtt{obs}_u$ for the function $\mathtt{obs}(u, \cdot)$, and $\mathtt{obs}_u(\alpha)$ for $\mathtt{obs}_u(s_0 \cdot \alpha)$, where $\alpha \in A^*$.

Given a sequence $\alpha \in A^*$, the *view* of a group of domains G of α is the sequence of the group's observations and the actions that belong to members of the group. Intuitively, G's view is the history of its observations and the actions it has performed. The function $\mathtt{view}_G$ defines the view of domain G. We first define the observation of group G at state s by $\mathtt{obs}_G(s) = \langle \mathtt{obs}_u(s)\rangle_{u \in G}$, i.e. the tuple of observations of individuals $u \in G$. The view function is then defined inductively by

$$\mathtt{view}_G(\epsilon) = \mathtt{obs}_G(s_0)$$
$$\mathtt{view}_G(\alpha a) = \begin{cases} \mathtt{view}_G(\alpha)\, a\, \mathtt{obs}_G(\alpha a) & \text{if } \mathtt{dom}(a) \in G \\ \mathtt{view}_G(\alpha) \circ \mathtt{obs}_G(\alpha a) & \text{otherwise} \end{cases}$$

for $\alpha \in A^*$ and $a \in A$. To capture that the semantics is asynchronous and insensitive to stuttering of observations, the definition uses the absorptive concatenation operator $\circ$: for set X, sequence $\alpha \in X^*$, and element $x \in X$, $\alpha \circ x = \alpha$

if x is equal to the last element of α, and $\alpha \circ x = \alpha x$ otherwise. When $G = \{u\}$ is a singleton, we write $\mathtt{view}_u$ for $\mathtt{view}_G$.

Finally, for any sequence of actions $\alpha \in A^*$, we write $\alpha \upharpoonright G$ for the subsequence of α of actions whose domain is in the set G.

2.4 Semantics

A machine satisfies an architecture if, in all possible executions of the machine, information flow is in accordance with the architecture's information-flow policy. We formalize this using an approach proposed by van der Meyden [2007], which involves the use of a concrete operational model to define an upper bound on the information that a domain is permitted to learn.

The operational model is captured using a function $\mathtt{ta}_u$, which maps a sequence of actions $\alpha \in A^*$ to a representation of the maximal information that domain u is permitted to have after α, according to the policy $\rightarrowtail$. (Term "ta" is derived from *transmission* of information about *actions*; the definition corrects problems identified by van der Meyden [2007] with earlier "intransitive purge" based semantics [Rushby, 1992].)

An action of v should convey information to u only if $v \rightarrowtail u$. Moreover, the information conveyed should be no more than the information that v is permitted to have. Given machine $M = \langle S, s_0, A, D, \mathtt{step}, \mathtt{obs}, \mathtt{dom}\rangle$, function $\mathtt{ta}_u$ is defined inductively by $\mathtt{ta}_u(\epsilon) = \epsilon$, and, for $\alpha \in A^*$ and $a \in A$,

$$\mathtt{ta}_u(\alpha a) = \begin{cases} \mathtt{ta}_u(\alpha) & \text{if } \mathtt{dom}(a) \not\rightarrowtail u \\ (\mathtt{ta}_u(\alpha), \mathtt{ta}_{\mathtt{dom}(a)}(\alpha), a) & \text{otherwise .} \end{cases}$$

Note that if information is not allowed to flow from $\mathtt{dom}(a)$ to u, then $\mathtt{ta}_u(\alpha a) = \mathtt{ta}_u(\alpha)$, i.e., the maximal information permitted to u does not change. If information is allowed to flow from $\mathtt{dom}(a)$ to u, then the information conveyed is at most the information that domain $\mathtt{dom}(a)$ is permitted to have (i.e., $\mathtt{ta}_{\mathtt{dom}(a)}(\alpha)$), and the action a that was performed. Thus, in this case we add the tuple $(\mathtt{ta}_{\mathtt{dom}(a)}(\alpha), a)$ to the maximal information $\mathtt{ta}_u(\alpha)$ that u was permitted to have before the action a was performed.

A machine is *TA-compliant* with an architecture if it has an appropriate set of domains, and for each domain u, what u observes in state $s_0 \cdot \alpha$ is determined by $\mathtt{ta}_u(\alpha)$. That is, $\mathtt{ta}_u$ describes the maximal information that u may learn: if in two runs α and α' the maximal information that u may learn is identical ($\mathtt{ta}_u(\alpha) = \mathtt{ta}_u(\alpha')$), then u's observations in each run must be identical ($\mathtt{obs}_u(\alpha) = \mathtt{obs}_u(\alpha')$).

Definition 1 (TA-compliance) *A system M is TA-compliant with architecture $(D, \rightarrowtail)$ if it has domains D and for all $u \in D$ and all sequences $\alpha, \alpha' \in A^*$ such that $\mathtt{ta}_u(\alpha) = \mathtt{ta}_u(\alpha')$, we have $\mathtt{obs}_u(\alpha) = \mathtt{obs}_u(\alpha')$.*

TA-compliance requires that if $\mathtt{ta}_u(\alpha) = \mathtt{ta}_u(\alpha')$ then the observations of u in state $s_0 \cdot \alpha$ and in state $s_0 \cdot \alpha'$ are equal. The following lemma shows that in fact TA-compliance implies that if $\mathtt{ta}_u(\alpha) = \mathtt{ta}_u(\alpha')$ then $\mathtt{view}_u(\alpha) = \mathtt{view}_u(\alpha')$.

Lemma 1 *If M is TA-compliant with respect to architecture $(D, \rightarrowtail)$, then for all agents u and all $\alpha, \alpha' \in A^*$ such that $\mathtt{ta}_u(\alpha) = \mathtt{ta}_u(\alpha')$ we have $\mathtt{view}_u(\alpha) = \mathtt{view}_u(\alpha')$.*

Proof: By induction on $|\alpha| + |\alpha'|$. The base case is trivial. Suppose that the result holds for all sequences of shorter combined length, and consider strings αa and α' such that $\mathtt{ta}_u(\alpha a) = \mathtt{ta}_u(\alpha')$.

We consider several cases.

- If $\mathtt{dom}(a) \not\rightarrowtail u$ then $\mathtt{ta}_u(\alpha) = \mathtt{ta}_u(\alpha a) = \mathtt{ta}_u(\alpha')$, so $\mathtt{view}_u(\alpha) = \mathtt{view}_u(\alpha')$, by the induction hypothesis, and $\mathtt{obs}_u(\alpha a) = \mathtt{obs}_u(\alpha)$, by TA-compliance. By reflexivity of $\rightarrowtail$, we have $\mathtt{dom}(a) \neq u$, so it follows that

$$\begin{aligned}\mathtt{view}_u(\alpha a) &= \mathtt{view}_u(\alpha) \circ \mathtt{obs}_u(\alpha a)\\ &= \mathtt{view}_u(\alpha)\\ &= \mathtt{view}_u(\alpha')\end{aligned}$$

- If $\mathtt{dom}(a) \rightarrowtail u$ then $\mathtt{ta}_u(\alpha a) = (\mathtt{ta}_u(\alpha), \mathtt{ta}_{\mathtt{dom}(a)}(\alpha), a) = \mathtt{ta}_u(\alpha')$. This means that $\alpha' \neq \epsilon$. Assume $\alpha' = \beta b$. Without loss of generality, we may assume $\mathtt{dom}(b) \rightarrowtail u$, since otherwise we may apply the previous case with the roles of αa and βb reversed. Let $\mathtt{ta}_u(\beta b) = (\mathtt{ta}_u(\beta), \mathtt{ta}_{\mathtt{dom}(b)}(\beta), b)$. Since $\mathtt{ta}_u(\alpha a) = \mathtt{ta}_u(\beta b)$, we have $a = b$, and by TA-compliance we have $\mathtt{obs}_u(\alpha a) = \mathtt{obs}_u(\beta b)$. Also, $\mathtt{ta}_u(\alpha) = \mathtt{ta}_u(\beta)$, and so by the inductive hypothesis, $\mathtt{view}_u(\alpha) = \mathtt{view}_u(\beta)$.
  - If $\mathtt{dom}(a) \neq u$, then

$$\begin{aligned}\mathtt{view}_u(\alpha a) &= \mathtt{view}_u(\alpha) \circ \mathtt{obs}_u(\alpha a)\\ &= \mathtt{view}_u(\beta) \circ \mathtt{obs}_u(\beta b)\\ &= \mathtt{view}_u(\beta b).\end{aligned}$$

  - If $\mathtt{dom}(a) = u$ then

$$\begin{aligned}\mathtt{view}_u(\alpha a) &= \mathtt{view}_u(\alpha)\, a\, \mathtt{obs}_u(\alpha a)\\ &= \mathtt{view}_u(\beta)\, b\, \mathtt{obs}_u(\beta b)\\ &= \mathtt{view}_u(\beta b).\end{aligned}$$

□

3 Information security properties

We use a (fairly standard) propositional epistemic logic [Fagin et al., 1995b] to express information security properties. The syntax is defined as follows:

$$\phi, \psi ::= \top \mid p \mid \neg\phi \mid \phi \wedge \psi \mid K_G\phi$$

G ranges over groups of domains.

$$
\begin{array}{ll}
M, \pi, \alpha \vDash \top & \\
M, \pi, \alpha \vDash p & \text{iff } \alpha \in \pi(p) \\
M, \pi, \alpha \vDash \neg\phi & \text{iff } M, \pi, \alpha \nvDash \phi \\
M, \pi, \alpha \vDash \phi \wedge \psi & \text{iff } M, \pi, \alpha \vDash \phi \text{ and } M, \pi, \alpha \vDash \psi \\
M, \pi, \alpha \vDash K_G\phi & \text{iff } M, \pi, \alpha' \vDash \phi \text{ for all } \alpha' \in A^* \text{ s.t. } \alpha \approx_G \alpha'
\end{array}
$$

Figure 1: Epistemic Logic Semantics

In case $G = \{u\}$ is a singleton, we write simply $K_u\phi$ for $K_G\phi$.

Formulas $\top$, p, $\neg\phi$, and $\phi \wedge \psi$ are standard from propositional logic: $\top$ is always satisfied, and p is a propositional constant. Epistemic formula $K_G\phi$ says that the group of domains G, considered as a single domain, knows ϕ.

Formulas are interpreted using a possible worlds semantics, where a world is a sequence of actions $\alpha \in A^*$. A *proposition* is a set $X \subseteq A^*$. We say proposition X is *non-trivial* if $X \neq \emptyset$ and $X \neq A^*$. An *interpretation function* π is a function from propositional constants to propositions.

We define the semantics of the logic using satisfaction relation $M, \pi, \alpha \vDash \phi$, which intuitively means that formula ϕ is true given interpretation function π, and machine M that has executed sequence $\alpha \in A^*$. Figure 1 defines relation $M, \pi, \alpha \vDash \phi$. We write $M, \pi \vDash \phi$ if for all $\alpha \in A^*$ we have $M, \pi, \alpha \vDash \phi$. We say that ϕ is *valid* if $M, \pi \vDash \phi$ for all systems M and interpretations π.

To interpret epistemic formulas $K_G\phi$, we use an indistinguishability relation for each group of domains G, that describes what sequences of actions G considers possible given its view of the actual sequence of actions. Two sequences of actions $\alpha \in A^*$ and $\alpha' \in A^*$ are indistinguishable to group of domains G, written $\alpha \approx_G \alpha'$, if G's views of the two sequences are identical: $\alpha \approx_G \alpha' \iff \texttt{view}_G(\alpha) = \texttt{view}_G(\alpha')$.

3.1 Group knowledge vs. Distributed knowledge

We note that the notion of group knowledge K_G differs from *distributed knowledge* [Fagin et al., 1995a], the notion most commonly used in the literature on epistemic logic for the knowledge that a group would have if they pooled their local information. The distributed knowledge operator D_G for group G is given semantics by

$$M, \pi, \alpha \vDash D_G\phi \text{ iff } M, \pi, \alpha' \vDash \phi \text{ for all } \alpha' \in A^* \text{ s.t. } \alpha \approx_G^D \alpha'$$

using a different indistinguishability relation $\approx_G^D$, defined as the intersection of $\approx_u$ for $u \in G$. The reason we use group knowledge is that it proves to have a stronger relationship to a type of architectural abstraction that we consider below. The two notions are related by the following result.

Lemma 2 *For $u \in G$, the formulas $K_u\phi \Rightarrow D_G\phi$ and $D_G\phi \Rightarrow K_G\phi$ are valid.*

Proof: That $K_u\phi \Rightarrow D_G\phi$ is valid is a well-known fact of epistemic logic. It follows simply from the definitions by noting that if $M, \pi, \alpha \models \neg D_G\phi$ then there

exists $\alpha' \approx_G^D \alpha$ such that $M, \pi, \alpha' \models \neg\phi$. Since $\alpha' \approx_G^D \alpha$ implies $\alpha' \approx_u \alpha$, we conclude $M, \pi, \alpha \models \neg K_u\phi$.

For $D_G\phi \Rightarrow K_G\phi$, we similarly consider the contrapositive. If $M, \pi, \alpha \models \neg K_G\phi$ then there exists $\alpha' \approx_G \alpha$ such that $M, \pi, \alpha' \models \neg\phi$. We claim that $\alpha' \approx_G \alpha$ implies $\alpha' \approx_u \alpha$ for all $u \in G$. It then follows that $\alpha' \approx_G^D \alpha$, and hence $M, \pi, \alpha \models \neg D_G\phi$.

To prove the claim, we show that there exists a function h such that for any $\beta \in A^*$, $h(\mathtt{view}_G(\beta)) = \mathtt{view}_u(\beta)$. The domain of this function is the set of possible views of G in M. These are nonempty sequences of actions a with $\mathtt{dom}(a) \in G$ and tuples $\langle \mathtt{obs}_u(s)\rangle_{u\in G}$ of observations. The first and last element of such a sequence is a tuple of observations, and the sequence does not contain any adjacent actions. We define the function h on such sequences as follows, where δ ranges over sequences produced by $\mathtt{view}_G$, $a \in A$, and t is a tuple of observations indexed by G. We write t_u for the component of t corresponding to $u \in G$. (In particular, note that if $t = \mathtt{obs}_G(s)$, then $t_u = \mathtt{obs}_u(s)$.)

$$
\begin{aligned}
h(t) &= t_u \\
h(\delta\, a\, t) &= \begin{cases} h(\delta)\, a\, t_u & \text{if } \mathtt{dom}(a) = u \\ h(\delta) \circ t_u & \text{otherwise} \end{cases} \\
h(\delta\, t) &= h(\delta) \circ t_u
\end{aligned}
$$

We prove that $h(\mathtt{view}_G(\beta)) = \mathtt{view}_u(\beta)$ by induction on $|\beta|$. The base case is trivial. Suppose that $h(\mathtt{view}_G(\beta)) = \mathtt{view}_u(\beta)$ and consider βb.

Suppose first that $\mathtt{dom}(b) \notin G$, so also $\mathtt{dom}(b) \neq u$. Then, $\mathtt{view}_G(\beta b) = \mathtt{view}_G(\beta) \circ \mathtt{obs}_G(\beta b)$.

- If $\mathtt{obs}_G(\beta b) = \mathtt{obs}_G(\beta)$, then $\mathtt{view}_G(\beta b)) = \mathtt{view}_G(\beta)$, and also $\mathtt{obs}_u(\beta b) = \mathtt{obs}_u(\beta)$. Moreover, $\mathtt{view}_u(\beta b) = \mathtt{view}_u(\beta) \circ \mathtt{obs}_u(\beta b) = \mathtt{view}_u(\beta)$. Thus, using the induction hypothesis, we have $h(\mathtt{view}_G(\beta b)) = h(\mathtt{view}_G(\beta)) = \mathtt{view}_u(\beta) = \mathtt{view}_u(\beta b)$.

- Otherwise, if $\mathtt{obs}_G(\beta b) \neq \mathtt{obs}_G(\beta)$, then $\mathtt{view}_G(\beta b) = \mathtt{view}_G(\beta)\, \mathtt{obs}_G(\beta b)$, so $h(\mathtt{view}_G(\beta b)) = h(\mathtt{view}_G(\beta)) \circ \mathtt{obs}_u(\beta b) = \mathtt{view}_u(\beta) \circ \mathtt{obs}_u(\beta b) = \mathtt{view}_u(\beta b)$.

Alternately, suppose $\mathtt{dom}(b) \in G$. Then $\mathtt{view}_G(\beta b) = \mathtt{view}_G(\beta)\, b\, \mathtt{obs}_G(\beta b)$. If $\mathtt{dom}(b) = u$ then $h(\mathtt{view}_G(\beta b)) = h(\mathtt{view}_G(\beta))\, b\, \mathtt{obs}_u(\beta b)$. If $\mathtt{dom}(b) \neq u$ then $h(\mathtt{view}_G(\beta b)) = h(\mathtt{view}_G(\beta)) \circ \mathtt{obs}_u(\beta b)$. Either way, by the inductive hypothesis, we have $h(\mathtt{view}_G(\beta b)) = \mathtt{view}_u(\beta b)$.

Thus we have $h(\mathtt{view}_G(\beta)) = \mathtt{view}_u(\beta)$ for any sequence β. Let $\alpha, \alpha' \in A^*$, and assume $\mathtt{view}_G(\alpha) = \mathtt{view}_G(\alpha')$. Then

$$
\begin{aligned}
\mathtt{view}_u(\alpha) &= h(\mathtt{view}_G(\alpha)) \\
&= h(\mathtt{view}_G(\alpha')) \\
&= \mathtt{view}_u(\alpha').
\end{aligned}
$$

□

The converse relationship $K_G\phi \Rightarrow D_G\phi$ is not valid. For example, consider a system M with exactly two domains u, v, which both make observation $\perp$ at all states. Let $G = \{u, v\}$. Consider the proposition p with $\pi(p)$ consisting of all sequences in which there is an action of domain u that precedes any action of domain v. Let $\alpha = a_u a_v$ and $\alpha' = a_v a_u$ where a_u is an action of u and a_v is an action of v. Then $M, \pi, \alpha \models K_G p$, since $\mathtt{view}_G(\alpha) = \mathtt{view}_G(\beta)$ implies that $\beta = \alpha$. However, we have $\mathtt{view}_u(\alpha) = \perp a_u \perp = \mathtt{view}_u(\alpha')$, and similarly for domain v, so $\alpha \approx_G^D \alpha'$. Since $M, \pi, \alpha' \not\models p$, we obtain that $M, \pi, \alpha \not\models D_G p$.[1]

3.2 Knowledge and architectural refinement

The value of using the less common notion $K_G\phi$ of group knowledge rather than distributed knowledge $D_G\phi$ is that it captures the way that knowledge properties are preserved under a particular type of architectural abstraction. Given a system $M = \langle S, s_0, A, D_1, \mathtt{step}, \mathtt{obs}, \mathtt{dom}\rangle$ and a surjective mapping $r : D_1 \to D_2$, define $r(M) = \langle S, s_0, A, D_2, \mathtt{step}, \mathtt{obs}', \mathtt{dom}'\rangle$ to be the system that is identical to M, except that it has domains D_2, and the functions $\mathtt{dom}'$ and $\mathtt{obs}'$ are defined by $\mathtt{dom}' = r \circ \mathtt{dom}$, and, for $u \in D_2$, $\mathtt{obs}'_u(s) = \mathtt{obs}_G(s)$, where $G = r^{-1}(u)$. Intuitively, each domain u in $r(M)$ corresponds to the group of domains $r^{-1}(u)$ in M, with every action of a domain in $r^{-1}(u)$ treated as an action of u. Similarly, for a formula ϕ of the epistemic logic, we write $r^{-1}(\phi)$ for the formula obtained by replacing each occurrence of a group G in a modal operator in ϕ by the group $r^{-1}(G)$.

The existence of a surjective mapping $r : D_1 \to D_2$ is one requirement for architectural refinement, which we discuss in greater detail in later sections. The following result shows that abstracting a system M to $r(M)$ (or, conversely, refining $r(M)$ to M) preserves satisfaction of formulas, subject to a corresponding abstraction on groups being applied in the formulas.

Theorem 1 *Let $r : D_1 \to D_2$ be surjective and let M be a system with domains D_1. Then for all interpretations π, sequences of actions α of M, and formulas ϕ (not including distributed knowledge operators D_G) for agents D_2 we have $r(M), \pi, \alpha \models \phi$ iff $M, \pi, \alpha \models r^{-1}(\phi)$.*

Proof: We first claim that for all $\alpha, \beta \in A^*$ and groups $G \subseteq D_2$, we have $\mathtt{view}^M_{r^{-1}(G)}(\alpha) = \mathtt{view}^M_{r^{-1}(G)}(\beta)$ iff $\mathtt{view}^{r(M)}_G(\alpha) = \mathtt{view}^{r(M)}_G(\beta)$, where the superscripts indicate the system within which views are computed. Note that, by definition,

$$\mathtt{obs}^{r(M)}_G(s) = \langle \mathtt{obs}^{r(M)}_u(s)\rangle_{u\in G} = \langle\langle \mathtt{obs}^M_v(s)\rangle_{v\in r^{-1}(u)}\rangle_{u\in G}$$

and

$$\mathtt{obs}^M_{r^{-1}(G)}(s) = \langle \mathtt{obs}^M_v(s)\rangle_{v\in r^{-1}(G)} = \langle \mathtt{obs}^M_v(s)\rangle_{v\in r^{-1}(u),\ u\in G}.$$

[1] We remark that the example relies upon the assumption of asynchrony: it can be shown that in synchronous systems we have $K_G\phi \equiv D_G\phi$.

Since these expressions simply group the same collection of values indexed by v in two different ways, there exists functions f and f^{-1} such that for all states s, we have $f(\mathtt{obs}^M_{r^{-1}(G)}(s)) = \mathtt{obs}^{r(M)}_G(s)$ and $f^{-1}(\mathtt{obs}^{r(M)}_G(s)) = \mathtt{obs}^M_{r^{-1}(G)}(s)$. Moreover, we have for states s and t that $\mathtt{obs}^{r(M)}_G(s) = \mathtt{obs}^{r(M)}_G(t)$ iff $\mathtt{obs}^M_{r^{-1}(G)}(s) = \mathtt{obs}^M_{r^{-1}(G)}(t)$.

Given a function h on (group) observations, we generalize it to a function h_+ on view-like sequences of actions and observations, inductively by $h_+(o) = h(o)$, for o an observation, $h_+(\sigma a o) = h_+(\sigma)a h(o)$ for σ a sequence, a an action and o an observation, and $h_+(\sigma o_1 o_2) = h_+(\sigma o_1) \circ h(o_2)$ for σ a sequence and o_1, o_2 observations.

We now prove the claim by showing that for all sequences of actions α, we have $f_+(\mathtt{view}^M_{r^{-1}(G)}(\alpha)) = \mathtt{view}^{r(M)}_G(\alpha))$ and $f_+^{-1}(\mathtt{view}^{r(M)}_G(\alpha)) = \mathtt{view}^M_{r^{-1}(G)}(\alpha)$. The proof is by induction on α. We consider just the case of f_+. In case $\alpha = \epsilon$, we have $f_+(\mathtt{view}^M_{r^{-1}(G)}(\alpha)) = f_+(\mathtt{obs}^M_{r^{-1}(G)}(s_0)) = f(\mathtt{obs}^M_{r^{-1}(G)}(s_0)) = \mathtt{obs}^{r(M)}_G(s_0) = \mathtt{view}^{r(M)}_G(\alpha)$. For sequences αa, there are two cases. If $\mathtt{dom}^M(a) \in r^{-1}(G)$ then $\mathtt{dom}^{r(M)}(a) = r(\mathtt{dom}^M(a)) \in G$. Thus, in this case, using the inductive hypothesis,

$$\begin{aligned}
f_+(\mathtt{view}^M_{r^{-1}(G)}(\alpha a)) &= f_+(\mathtt{view}^M_{r^{-1}(G)}(\alpha)\, a\, \mathtt{obs}^M_{r^{-1}(G)}(s_0 \cdot \alpha a)) \\
&= f_+(\mathtt{view}^M_{r^{-1}(G)}(\alpha))\, a\, f(\mathtt{obs}^M_{r^{-1}(G)}(s_0 \cdot \alpha a)) \\
&= \mathtt{view}^{r(M)}_G(\alpha)\, a\, \mathtt{obs}^{r(M)}_G(s_0 \cdot \alpha a) \\
&= \mathtt{view}^{r(M)}_G(\alpha a).
\end{aligned}$$

Alternately, if $\mathtt{dom}^M(a) \notin r^{-1}(G)$ then $\mathtt{dom}^{r(M)}(a) = r(\mathtt{dom}^M(a)) \notin G$. In this case, there are two further possibilities. If $\mathtt{obs}^M_{r^{-1}(G)}(s_0 \cdot \alpha) = \mathtt{obs}^M_{r^{-1}(G)}(s_0 \cdot \alpha a)$, then also $\mathtt{obs}^{r(M)}_G(s_0 \cdot \alpha) = \mathtt{obs}^{r(M)}_G(s_0 \cdot \alpha a)$. Thus

$$\begin{aligned}
f_+(\mathtt{view}^M_{r^{-1}(G)}(\alpha a)) &= f_+(\mathtt{view}^M_{r^{-1}(G)}(\alpha) \circ \mathtt{obs}^M_{r^{-1}(G)}(s_0 \cdot \alpha a)) \\
&= f_+(\mathtt{view}^M_{r^{-1}(G)}(\alpha)) \\
&= \mathtt{view}^{r(M)}_G(\alpha) \\
&= \mathtt{view}^{r(M)}_G(\alpha) \circ \mathtt{obs}^{r(M)}_G(s_0 \cdot \alpha a) \\
&= \mathtt{view}^{r(M)}_G(\alpha a).
\end{aligned}$$

The other possibility is that $\mathtt{obs}^M_{r^{-1}(G)}(s_0 \cdot \alpha) \neq \mathtt{obs}^M_{r^{-1}(G)}(s_0 \cdot \alpha a)$, where we have

$$\begin{aligned}
f_+(\mathtt{view}^M_{r^{-1}(G)}(\alpha a)) &= f_+(\mathtt{view}^M_{r^{-1}(G)}(\alpha)\, \mathtt{obs}^M_{r^{-1}(G)}(s_0 \cdot \alpha a)) \\
&= f_+(\mathtt{view}^M_{r^{-1}(G)}(\alpha)) \circ f(\mathtt{obs}^M_{r^{-1}(G)}(s_0 \cdot \alpha a)) \\
&= \mathtt{view}^{r(M)}_G(\alpha) \circ \mathtt{obs}^{r(M)}_G(s_0 \cdot \alpha a) \\
&= \mathtt{view}^{r(M)}_G(\alpha a).
\end{aligned}$$

Thus, in any case we have $f_+(\mathtt{view}^M_{r^{-1}(G)}(\alpha a)) = \mathtt{view}^{r(M)}_G(\alpha a)$, completing the induction. The argument for f_+^{-1} is symmetric.

The result now follows by induction on the construction of ϕ. The cases of atomic propositions and boolean operators are trivial. For formulas of the form $K_G\psi$, we have $r(M), \pi, \alpha \models K_G(\psi)$ iff $r(M), \pi, \beta \models \psi$ for all sequences of actions β with $\mathtt{view}^{r(M)}_G(\alpha) = \mathtt{view}^{r(M)}_G(\beta)$. By the induction hypothesis and the claim proved above, this is equivalent to $M, \pi, \beta \models r^{-1}(\psi)$ for all sequences of actions β with $\mathtt{view}^M_{r^{-1}(G)}(\alpha) = \mathtt{view}^M_{r^{-1}(G)}(\beta)$. This is equivalent to $M, \pi, \alpha \models K_{r^{-1}(G)} r^{-1}(\psi)$, i.e., $M, \pi, \alpha \models r^{-1}(K_G\psi)$. □

We note that the example given above to show the difference between group and distributed knowledge also shows that Theorem 1 would not hold if we were to include distributed knowledge in the language and analogously define $r^{-1}(D_G\phi) = D_{r^{-1}(G)}\phi$. For example, consider the function r with $r(u) = r(v) = w$. Since $r^{-1}(K_w p) = K_{\{u,v\}}p$ and $M, \pi, \alpha \models K_{\{u,v\}}p$, we have $r(M), \pi, \alpha \models K_w p$, hence $r(M), \pi, \alpha \models D_w p$. However, as shown above, we do not have $M, \pi, \alpha \models D_{\{u,v\}}p$, i.e., we do not have $M, \pi, \alpha \models r^{-1}(D_w p)$.

G-dependent propositions A proposition in a system M depends on the actions of a group G if its truth value can be affected by making changes only to the actions of domains in the group G. The notion of dependence of a proposition on a group G is useful to specify confidential information in some of our examples.

Formally, for a group G of domains and $\alpha \in A^*$, say that a proposition $X \subseteq A^*$ *depends on G actions at α* if there exists $\beta \in A^*$ such that $\alpha \restriction \overline{G} = \beta \restriction \overline{G}$ but $\alpha \in X$ iff $\beta \notin X$. (Notation $\overline{G}$ is shorthand for the set $D \setminus G$, and denotes the set of all domains excluding those in G.) Intuitively, this says that which G actions have occurred, and their placement with respect to the actions of other domains, can affect whether or not the proposition holds. We say that X *depends everywhere on G actions* if X depends on G actions at α for all $\alpha \in A^*$.[2]

We can also reason about how architectural abstraction affects G-dependent propositions.

Lemma 3 *Let $r : D_1 \to D_2$ be surjective and let M be a system with domains D_1. Then proposition X depends on G actions at α in $r(M)$ iff X depends on $r^{-1}(G)$ actions at α in M.*

Proof: Note that the restriction operation is relative to a system M; to emphasize this we write $\alpha \restriction^M G$ to indicate that we use the domain function from M. Note also that

$$r^{-1}(\overline{G}) = \overline{r^{-1}(G)} \ .$$

[2]In earlier versions of this work, we used nontrivial G-action local propositions. A proposition X is *G-action local* if for all $\alpha, \beta \in A^*$, if $\alpha \restriction G = \beta \restriction G$, then $\alpha \in X \iff \beta \in X$. It can easily be seen that a non-trivial G-action local proposition is everywhere dependent on G actions. Consequently, a formulation of our results using propositions that depend on G actions is more general.

By definition, X depends on G actions at α in $r(M)$ if there exists $\beta \in A^*$ such that $\alpha \restriction^{r(M)} \overline{G} = \beta \restriction^{r(M)} \overline{G}$ and $\alpha \in X$ iff $\beta \notin X$. Similarly, X depends on $r^{-1}(G)$ actions at α in M if there exists $\beta \in A^*$ such that $\alpha \restriction^{M} \overline{r^{-1}(G)} = \beta \restriction^{r(M)} \overline{r^{-1}(G)}$ and $\alpha \in X$ iff $\beta \notin X$.

Since $\mathtt{dom}^M(a) \in \overline{r^{-1}(G)}$ iff $\mathtt{dom}^M(a) \in r^{-1}(\overline{G})$ iff $\mathtt{dom}^{r(M)}(a) \in \overline{G}$, we have

$$\alpha \restriction^{r(M)} \overline{G} = \beta \restriction^{r(M)} \overline{G} \text{ iff } \alpha \restriction^{M} \overline{r^{-1}(G)} = \beta \restriction^{M} \overline{r^{-1}(G)} \ .$$

It follows that X depends on G actions at α in $r(M)$ iff X depends on $r^{-1}(G)$ actions at α in M. □

3.3 Example: $\mathcal{HL}$ information security

The logic allows us to state information security properties about machines, in terms of the knowledge of domains. The architecture can provide sufficient structure to prove that a given information security property holds in all machines that comply with the architecture.

For example, using the $\mathcal{HL}$ architecture, we are able to show that in any execution of any machine that complies with $\mathcal{HL}$, the domain L does not know any proposition that depends on H actions.

Theorem 2 *If M is TA-compliant with $\mathcal{HL}$ and $\pi(p)$ depends on H actions at α then $M, \pi, \alpha \vDash \neg K_L p$.*

Proof: We first show, for any sequence $\gamma \in A^*$, that $\mathtt{ta}_L(\gamma) = \mathtt{ta}_L(\gamma \restriction \{L\})$. The base case, $\gamma = \epsilon$ is trivial. Consider γa. If $\mathtt{dom}(a) \neq L$ then

$$\begin{aligned} \mathtt{ta}_L(\gamma a) &= \mathtt{ta}_L(\gamma) && \text{by defn } \mathtt{ta}_L \\ &= \mathtt{ta}_L(\gamma \restriction \{L\}) && \text{by IH} \\ &= \mathtt{ta}_L(\gamma a \restriction \{L\}) && \text{since } \mathtt{dom}(a) \neq L. \end{aligned}$$

If $\mathtt{dom}(a) = L$ then

$$\begin{aligned} \mathtt{ta}_L(\gamma a) &= (\mathtt{ta}_L(\gamma), \mathtt{ta}_L(\gamma), a) && \text{by defn } \mathtt{ta}_L \\ &= (\mathtt{ta}_L(\gamma \restriction \{L\}), \mathtt{ta}_L(\gamma \restriction \{L\}), a) && \text{by IH} \\ &= \mathtt{ta}_L(\gamma a \restriction \{L\}) && \text{since } \mathtt{dom}(a) = L. \end{aligned}$$

Now, since $\pi(p)$ depends on H actions at α there exists $\beta \in A^*$ such that $\alpha \restriction \{L\} = \beta \restriction \{L\}$ and $\alpha \in \pi(p)$ iff $\beta \notin \pi(p)$. By what was shown above, $\mathtt{ta}_L(\alpha) = \mathtt{ta}_L(\alpha \restriction \{L\}) = \mathtt{ta}_L(\beta \restriction \{L\}) = \mathtt{ta}_L(\beta)$, so since M is TA-compliant, by Lemma 1, we have $\alpha \approx_L \beta$ and the result follows immediately. □

3.4 Example: Hinke-Schaefer

In the Hinke-Schaefer database architecture $\mathcal{HS}$, none of the domains L_{user}, L_{DBMS}, or L_F know anything about the domains H_{user}, H_{DBMS} or H_F. This is true even if we consider the group knowledge of L_{user}, L_{DBMS}, and L_F.

Theorem 3 *Let system M be TA-compliant with $\mathcal{HS}$, and let $G = \{L_{user}, L_{DBMS}, L_F\}$. If $\pi(p)$ depends on $\{H_{user}, H_{DBMS}, H_F\}$ actions at α, then $M, \pi, \alpha \vDash \neg K_G p$.*

Although we could prove Theorem 3 directly, we defer the proof to Section 5.3, where the result follows easily from Theorem 2 and the relationship between architectures $\mathcal{HL}$ and $\mathcal{HS}$. (Specifically, architecture $\mathcal{HS}$ is a refinement of architecture $\mathcal{HL}$.)

Since $K_u p \Rightarrow K_G p$ is valid for $u \in G$, it follows that also $M, \pi \vDash \neg K_u p$ for $u \in \{L_{user}, L_{DBMS}, L_F\}$. In particular, L_{user} does not have any information about the High side of the system.

4 Extended architecture

The architectures used so far impose coarse, global constraints on the causal structure of systems. If $u \rightarrowtail v$ then TA-compliance permits domain u to send to domain v any and all data it has. However, in many systems, key security properties depend on the fact that trusted components allow only certain information to flow from one domain to another. Finer specification of information flows in the architecture allow us to prove stronger information security properties.

In this section we extend the notion of architecture by introducing *filter functions* to allow fine-grained specification of what information flows between domains. We define semantics for these extended architectures, and present examples where the extended architectures allow us to prove strong information security properties.

4.1 Filter functions

An extended architecture is a pair $\mathcal{A} = (D, \rightarrowtail)$, where D is a set of security domains, and $\rightarrowtail \subseteq D \times D \times (\mathcal{L} \cup \{\top\})$, where $\mathcal{L}$ is a set of function names. We write $u \overset{f}{\rightarrowtail} v$ when $(u, v, f) \in \rightarrowtail$, write $u \rightarrowtail v$ as shorthand for $\exists f.\ (u, v, f) \in \rightarrowtail$, and $u \not\rightarrowtail v$ as shorthand for $\neg\exists f.\ (u, v, f) \in \rightarrowtail$.

Intuitively, $u \overset{f}{\rightarrowtail} v$ represents that information flow from u to v is permitted, but may be subject to constraints. In case $f = \top$, there are no constraints on information flow from u to v: any information that may be possessed by u is permitted to be passed to v when u acts, just as in the definition of TA-compliance. If $f \in \mathcal{L}$ then information is allowed to flow from domain u to domain v, but it needs to be *filtered* through the function denoted by f: only information output by this function may be transmitted from u to v. If $u \not\rightarrowtail v$ then no direct flow of information from u to v is permitted.

In some cases, it may be possible for the operating system or network infrastructure to enforce a given filter function. However, in general, a filter function is a local constraint on a trusted component of the system. That is, if $u \overset{f}{\rightarrowtail} v$ for $f \neq \top$, then component u is trusted to enforce that information sent to v is filtered appropriately.

We require that extended architectures have the following properties:

1. For all $u, v \in D$, there exists at most one $f \in \mathcal{L} \cup \{\top\}$ such that $u \overset{f}{\rightarrowtail} v$.
2. The relation $\rightarrowtail$ is reflexive in that for all $u \in D$ we have $(u, u, \top) \in \rightarrowtail$.

The first condition requires that all permitted flows of information from u to v are represented using a single labeled edge. Intuitively, any policy with multiple such edges can always be transformed into one satisfying this condition, by combining the pieces of information flowing across these edges into a tuple that flows across a single edge. The second condition is motivated from the fact, already noted above, that information flow from a domain to itself cannot be prevented.

For example, the following diagram shows an extended architecture with domains H, D, and L, intended to represent a high-security domain, a trusted declassifier, and a low-security domain. The arrows indicate permitted flow between domains. The label on the edge from D to L indicates that information going from D to L should be filtered by function *rel*. When drawing extended architectures, we annotate arrows between domains with the filter function names. For arrows drawn without a label, and elided reflexive arrows, the implied label is $\top$.

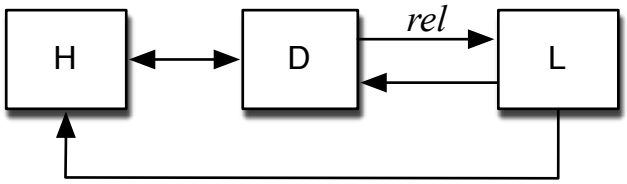

Extended architectures do not define the interpretations of the function names $\mathcal{L}$. If $\mathcal{A} = (D, \rightarrowtail)$ is an extended architecture, an *interpretation* for $\mathcal{A}$ is a tuple $\mathcal{I} = (A, \mathtt{dom}, \mathtt{I})$, where A is a set of actions, $\mathtt{dom} : A \to D$ assigns these actions to domains of $\mathcal{A}$, and $\mathtt{I}$ is a function mapping each $f \in \mathcal{L}$ to a function with domain $A^* \times A$ (and arbitrary codomain). We call the pair $(\mathcal{A}, \mathcal{I})$ an *interpreted extended architecture*, or simply an *interpreted architecture*.

Intuitively, if $u \overset{f}{\rightarrowtail} v$ and $\alpha \in A^*$ and $a \in A$ is an action with $\mathtt{dom}(a) = u$, then $\mathtt{I}(f)(\alpha, a)$ is the information that is permitted to flow from u to v when the action a is performed after occurrence of the sequence of actions $\alpha \in A^*$.

Given extended architecture $(D, \rightarrowtail)$ and an architectural interpretation $\mathcal{I} = (A, \mathtt{dom}, \mathtt{I})$, we define a function $\mathtt{fta}_u$ with domain A^* that, like $\mathtt{ta}_u$, captures the maximal information that domain u is permitted to have after a sequence of actions has been executed. The definition is recursive with a function $\mathtt{T}_{v,u}$ for $u, v \in D$, mapping a sequence $\alpha \in A^*$ and an action $a \in A$ with $\mathtt{dom}(a) = u$ to

$$\mathtt{T}_{v,u}(\alpha, a) = \begin{cases} \epsilon & \text{if } v \not\rightarrowtail u \\ (\mathtt{fta}_v(\alpha), a) & \text{if } v \overset{\top}{\rightarrowtail} u \\ \mathtt{I}(f)(\alpha, a) & \text{if } v \overset{f}{\rightarrowtail} u. \end{cases}$$

Intuitively, $\mathtt{T}_{v,u}(\alpha, a)$ represents the new information permitted to be known by u when action a is performed after sequence α.

The function $\mathtt{fta}_u$ is defined by $\mathtt{fta}_u(\epsilon) = \epsilon$, and, for $\alpha \in A^*$ and $a \in A$, $\mathtt{fta}_u(\alpha a) = \mathtt{fta}_u(\alpha)\,\hat{}\,\mathtt{T}_{\mathtt{dom}(a),u}(\alpha, a)$ where $\hat{}$ is the operation of appending an element to the end of a sequence. Some important technical points concerning the append operation are that for any sequence σ, we define $\sigma\,\hat{}\,\epsilon = \sigma$ (i.e., appending the empty sequence ϵ has no effect), and if δ happens to be a nonempty sequence, then $\sigma\,\hat{}\,\delta$ is the sequence that extends the sequence σ by the single additional element δ. For example if δ is the sequence ab, then $\sigma\,\hat{}\,\delta$ has final element equal to the sequence ab rather than b.

Unfolding the definition, we obtain

$$\mathtt{fta}_u(\alpha a) = \begin{cases} \mathtt{fta}_u(\alpha) & \text{if } \mathtt{dom}(a) \not\rightarrowtail u \\ \mathtt{fta}_u(\alpha)\,\hat{}\,(\mathtt{fta}_{\mathtt{dom}(a)}(\alpha), a) & \text{if } \mathtt{dom}(a) \overset{\top}{\rightarrowtail} u \\ \mathtt{fta}_u(\alpha)\,\hat{}\,\mathtt{I}(f)(\alpha, a) & \text{if } \mathtt{dom}(a) \overset{f}{\rightarrowtail} u. \end{cases}$$

The first two clauses resemble the definition of $\mathtt{ta}_u$; the third adds to this that the information flowing along an edge labeled by a function name f is filtered by the interpretation $\mathtt{I}(f)$. Note that if $\mathtt{I}(f)(\alpha, a) = \epsilon$, where $\mathtt{dom}(a) \overset{f}{\rightarrowtail} u$, then $\mathtt{fta}_u(\alpha a) = \mathtt{fta}_u(\alpha)$. That is, filter function $\mathtt{I}(f)$ can specify that no information should flow under certain conditions. Note also that $\mathtt{fta}_u$ and $\mathtt{T}_{v,u}$ have implicit parameters, viz., an information flow policy $\rightarrowtail$ and an architectural interpretation $\mathcal{I} = (A, \mathtt{dom}, \mathtt{I})$. When we need to make some of these parameters explicit, we write expressions such as $\mathtt{fta}_u^{(\rightarrowtail, \mathcal{I})}$, or $\mathtt{fta}_u^{\rightarrowtail}$.

The function $\mathtt{fta}_u$ is used analogously to $\mathtt{ta}_u$ to define the maximal information that a domain is permitted to observe for a given sequence of actions. However, $\mathtt{fta}_u$ is a more precise bound than $\mathtt{ta}_u$, as it uses filter functions to bound the information sent between domains.

It is reasonable to assume that information sent from u to v is information that u is permitted to have. We say that a function is $\mathtt{fta}_u$-compatible when the information it conveys is determined by information that u is permitted to have.

Definition 2 *Function h with domain $A^* \times A$ is $\mathtt{fta}_u$-compatible when for all sequences $\alpha, \beta \in A^*$, $\mathtt{fta}_u(\alpha) = \mathtt{fta}_u(\beta)$ implies that for all $a \in A$ with $\mathtt{dom}(a) = u$ we have $h(\alpha, a) = h(\beta, a)$.*

We say that the interpretation $\mathcal{I} = (A, \mathtt{dom}, \mathtt{I})$ is *compatible* with $\mathcal{A} = (D, \rightarrowtail)$ if for all $u \in D$ and edges $u \overset{f}{\rightarrowtail} v$ with $f \in \mathcal{L}$, the function $\mathtt{I}(f)$ is $\mathtt{fta}_u$-compatible. In what follows, we require that interpretations be compatible with their architectures.

A machine complies with an interpreted extended architecture if it has appropriate domains and actions, and for each domain u, what u observes in state $s_0 \cdot \alpha$ is determined by $\mathtt{fta}_u(\alpha)$. We call such a machine *FTA-compliant.* ("FTA" is derived from *filtered transmission* of information about *actions.*)

Definition 3 (FTA-compliant) *A machine* $M = \langle S, s_0, A, D, \mathtt{step}, \mathtt{obs}, \mathtt{dom}\rangle$ *is FTA-compliant with an interpreted architecture* $(\mathcal{A}, \mathcal{I})$*, with* $\mathcal{A} = (D', \rightarrowtail)$ *and* $\mathcal{I} = (A', \mathtt{dom}', \mathtt{I})$*, if* $A = A'$*,* $D = D'$*,* $\mathtt{dom} = \mathtt{dom}'$ *and for all agents* $u \in D$ *and all* $\alpha, \alpha' \in A^*$*, if* $\mathtt{fta}_u(\alpha) = \mathtt{fta}_u(\alpha')$ *then* $\mathtt{obs}_u(\alpha) = \mathtt{obs}_u(\alpha')$*.*

For an interpreted architecture $\mathcal{AI}$, with actions A and domains D, if π is an interpretation with $\pi(p) \subseteq A^*$, we write $\mathcal{AI}, \pi, \alpha \models \phi$ if $M, \pi, \alpha \models \phi$ for all systems M that are FTA-compliant with $\mathcal{AI}$. Similarly, we write $\mathcal{AI}, \pi \models \phi$ if $\mathcal{AI}, \pi, \alpha \models \phi$ for all $\alpha \in A^*$.

Separating extended architectures from their interpretations ensures that extended architectures can be completely represented by graphical diagrams with labeled edges. It also allows us to deal with examples where an extended architecture can be implemented in a variety of ways, and weak constraints on the set of actions and the set of filter functions suffice to enforce the security properties of interest. We will present a number of examples of this in what follows. To capture the constraints on the architectural interpretations at the semantic level, we use the notion of an *architectural specification*, which is a pair $(\mathcal{A}, \mathcal{C})$ where $\mathcal{A}$ is an extended architecture and $\mathcal{C}$ is a set of architectural interpretations for $\mathcal{A}$. (We will not attempt in this paper to develop any syntactic notation for architectural specifications.)

Definition 4 *A machine* M *is FTA-compliant with an architectural specification* $(\mathcal{A}, \mathcal{C})$ *if there exists an interpretation* $\mathcal{I} \in \mathcal{C}$ *such that* M *is FTA-compliant with the interpreted architecture* $(\mathcal{A}, \mathcal{I})$*.*

The following theorem shows that FTA-compliance generalizes TA-compliance. Thus, we are free to interpret a given architecture as an extended architecture.

Theorem 4 *Let* $\mathcal{A}_1 = (D, \rightarrowtail_1)$ *be an architecture, and let* $\mathcal{A}_2 = (D, \rightarrowtail_2)$ *be the extended architecture such that* $(u, v, f) \in \rightarrowtail_2$ *if and only if* $f = \top$ *and* $(u, v) \in \rightarrowtail_1$*. Let* M *be a machine with domains* D*, actions* A *and domain function* $\mathtt{dom}$*. Let* $\mathcal{I} = (A, \mathtt{dom}, \mathtt{I})$ *be any interpretation for* $\mathcal{A}_2$ *with this set of actions and domain function. Then* M *is TA-compliant with* $\mathcal{A}_1$ *if and only if* M *is FTA-compliant with* $(\mathcal{A}_2, \mathcal{I})$*.*

Proof: We show that for all $u \in D$, $\mathtt{ta}_u$ is isomorphic to $\mathtt{fta}_u$, in the sense that $\mathtt{ta}_u(\alpha) = \mathtt{ta}_u(\beta)$ iff $\mathtt{fta}_u(\alpha) = \mathtt{fta}_u(\beta)$. Let g be the function mapping nested triples to sequences, defined inductively by $g(\epsilon) = \epsilon$ and $g((t_1, t_2, a)) = g(t_1)\,\hat{}\,(g(t_2), a)$. Conversely, let h be the function from the range of g to nested tuples, defined inductively by $h(\epsilon) = \epsilon$ and $h(\sigma\,\hat{}\,(x, a)) = (h(\sigma), h(x), a))$. A straightforward induction shows that $h \circ g$ is the identity function, so g is 1-1. We show by induction on α that for all $u \in D$, $g(\mathtt{ta}_u(\alpha)) = \mathtt{fta}_u(\alpha)$.

The base case is trivial. Assume that for all $u \in D$ we have $g(\mathtt{ta}_u(\alpha)) = \mathtt{fta}_u(\alpha)$ and consider αa. Let $u \in D$. If $\mathtt{dom}(a) \not\rightarrowtail_1 u$ then $\mathtt{dom}(a) \not\rightarrowtail_2 u$ and

$$\begin{aligned} g(\mathtt{ta}_u(\alpha a)) &= g(\mathtt{ta}_u(\alpha)) && \text{by defn } \mathtt{ta}_u \\ &= \mathtt{fta}_u(\alpha) && \text{by IH} \\ &= \mathtt{fta}_u(\alpha a) && \text{by defn } \mathtt{fta}_u \end{aligned}$$

If $\mathtt{dom}(a) \rightarrowtail_1 u$ then $(\mathtt{dom}(a), u, \top) \in \rightarrowtail_2$, and

$$\begin{aligned} g(\mathtt{ta}_u(\alpha a)) &= g((\mathtt{ta}_u(\alpha), \mathtt{ta}_{\mathtt{dom}(a)}(\alpha), a)) && \text{by defn } \mathtt{ta}_u \\ &= g(\mathtt{ta}_u(\alpha)) \,\hat{}\, (g(\mathtt{ta}_{\mathtt{dom}(a)}(\alpha)), a) && \text{by defn } g \\ &= \mathtt{fta}_u(\alpha) \,\hat{}\, (\mathtt{fta}_{\mathtt{dom}(a)}(\alpha), a) && \text{by IH} \\ &= \mathtt{fta}_u(\alpha a) && \text{by defn } \mathtt{fta}_u \end{aligned}$$

It now follows that $\mathtt{ta}_u(\alpha) = \mathtt{ta}_u(\beta)$ iff $\mathtt{fta}_u(\alpha) = \mathtt{fta}_u(\beta)$. For, if $\mathtt{ta}_u(\alpha) = \mathtt{ta}_u(\beta)$ then $\mathtt{fta}_u(\alpha) = g(\mathtt{ta}_u(\alpha)) = g(\mathtt{ta}_u(\beta)) = \mathtt{fta}_u(\beta)$. Conversely, if $\mathtt{fta}_u(\alpha) = \mathtt{fta}_u(\beta)$ then $g(\mathtt{ta}_u(\alpha)) = g(\mathtt{ta}_u(\beta))$ and we get $\mathtt{ta}_u(\alpha) = \mathtt{ta}_u(\beta)$ by injectivity of g. □

FTA-compliance requires that if $\mathtt{fta}_u(\alpha) = \mathtt{fta}_u(\alpha')$ then the observations of u in state $s_0 \cdot \alpha$ and in state $s_0 \cdot \alpha'$ are equal. The following lemma (similar to Lemma 1) shows that in fact FTA-compliance implies that if $\mathtt{fta}_u(\alpha) = \mathtt{fta}_u(\alpha')$ then we have that $\mathtt{view}_u(\alpha) = \mathtt{view}_u(\alpha')$.

We require a technical assumption for this result: say that an interpreted architecture $(\mathcal{A}, \mathcal{I})$ is *non-conflating* if for all $u, v \in D$, if $u \rightarrowtail^{f} v$ with $f \neq \top$ then for all actions $a, b \in A$ with $\mathtt{dom}(a) = u$ and $\mathtt{dom}(b) = v$, and for all $\alpha, \beta \in A^*$ we have $\mathtt{I}(f)(\alpha, a) \neq (\mathtt{fta}_v(\beta), b)$. Recall that by reflexivity, $u \rightarrowtail^{f} v$ with $f \neq \top$ implies $u \neq v$. Intuitively, the condition states that, in the context of the definition of $\mathtt{fta}_v$, it is always possible for domain v to distinguish the type of information $\mathtt{I}(f)(\alpha, a)$ transmitted to it from the type of information $(\mathtt{fta}_v(\beta), b)$ transmitted by v to itself. That is, v can distinguish between the effects of its own actions on $\mathtt{fta}_v$ and the effects of other domains' actions. Since, intuitively, v should be aware of its own actions, it is reasonable to expect that this is generally satisfied.

Lemma 4 *If M is FTA-compliant with non-conflating interpreted architecture $(\mathcal{A}, \mathcal{I})$, with $\mathcal{A} = (D, \rightarrowtail)$ and $\mathcal{I} = (A, \mathtt{dom}, \mathtt{I})$, then for all agents $u \in D$ and all $\alpha, \alpha' \in A^*$ such that $\mathtt{fta}_u(\alpha) = \mathtt{fta}_u(\alpha')$ we have $\mathtt{view}_u(\alpha) = \mathtt{view}_u(\alpha')$*

Proof: By induction on $|\alpha| + |\alpha'|$. The base case is trivial. Suppose that the result holds for all sequences of shorter combined length, and consider strings αa and α' such that $\mathtt{fta}_u(\alpha a) = \mathtt{fta}_u(\alpha')$. Note that by FTA-compliance this implies that $\mathtt{obs}_u(\alpha a) = \mathtt{obs}_u(\alpha')$. We consider several cases.

- If $\mathtt{dom}(a) \not\rightarrowtail u$ then $\mathtt{fta}_u(\alpha) = \mathtt{fta}_u(\alpha a) = \mathtt{fta}_u(\alpha')$, so $\mathtt{view}_u(\alpha) = \mathtt{view}_u(\alpha')$, by the induction hypothesis, and $\mathtt{obs}_u(\alpha a) = \mathtt{obs}_u(\alpha)$, by FTA-compliance. By reflexivity of $\rightarrowtail$, we have $\mathtt{dom}(a) \neq u$, so it follows that

$$\begin{aligned} \mathtt{view}_u(\alpha a) &= \mathtt{view}_u(\alpha) \circ \mathtt{obs}_u(\alpha a) \\ &= \mathtt{view}_u(\alpha) \\ &= \mathtt{view}_u(\alpha') \end{aligned}$$

- If $\mathtt{dom}(a) \rightarrowtail^{f} u$ and $f \neq \top$ and $\mathtt{I}(f)(\alpha, a) = \epsilon$ then

$$\mathtt{fta}_u(\alpha) = \mathtt{fta}_u(\alpha) \mathbin{\hat{}} \mathtt{I}(f)(\alpha, a) = \mathtt{fta}_u(\alpha a) = \mathtt{fta}_u(\alpha').$$

  Thus we have $\mathtt{view}_u(\alpha) = \mathtt{view}_u(\alpha')$, by the induction hypothesis, and $\mathtt{obs}_u(\alpha a) = \mathtt{obs}_u(\alpha)$, by FTA-compliance. By reflexivity of $\rightarrowtail$, we have $\mathtt{dom}(a) \neq u$. Therefore

$$\begin{aligned} \mathtt{view}_u(\alpha a) &= \mathtt{view}_u(\alpha) \circ \mathtt{obs}_u(\alpha a) \\ &= \mathtt{view}_u(\alpha) \\ &= \mathtt{view}_u(\alpha') \end{aligned}$$

- If $\mathtt{dom}(a) \rightarrowtail^{f} u$ and $f \neq \top$ and $\mathtt{I}(f)(\alpha, a) \neq \epsilon$ then $\mathtt{fta}_u(\alpha') = \mathtt{fta}_u(\alpha a) = \mathtt{fta}_u(\alpha) \mathbin{\hat{}} \mathtt{I}(f)(\alpha, a)$ yields that $\alpha' \neq \epsilon$. Assume $\alpha' = \beta b$.

  - If either $\mathtt{dom}(b) \not\rightarrowtail u$, or $\mathtt{dom}(b) \rightarrowtail^{g} u$ and $g \neq \top$ and $\mathtt{I}(g)(\beta, b) = \epsilon$, then we can swap the roles of αa and βb and apply the argument from the previous two cases.
  - If $\mathtt{dom}(b) \rightarrowtail^{g} u$ and $g \neq \top$ and $\mathtt{I}(g)(\beta, b) \neq \epsilon$ then

$$\mathtt{fta}_u(\alpha a) = \mathtt{fta}_u(\alpha) \mathbin{\hat{}} \mathtt{I}(f)(\alpha, a) = \mathtt{fta}_u(\beta) \mathbin{\hat{}} \mathtt{I}(g)(\beta, b) = \mathtt{fta}_u(\beta b),$$

    and so $\mathtt{fta}_u(\alpha) = \mathtt{fta}_u(\beta)$, since neither of the appended elements is ϵ. By the inductive hypothesis, $\mathtt{view}_u(\alpha) = \mathtt{view}_u(\beta)$. Since $f \neq \top$ and $g \neq \top$, by the reflexivity of $\rightarrowtail$, we have $\mathtt{dom}(a) \neq u$ and $\mathtt{dom}(b) \neq u$. Thus

$$\begin{aligned} \mathtt{view}_u(\alpha a) &= \mathtt{view}_u(\alpha) \circ \mathtt{obs}_u(\alpha a) \\ &= \mathtt{view}_u(\beta) \circ \mathtt{obs}_u(\beta b) \\ &= \mathtt{view}_u(\beta b). \end{aligned}$$

  - If $\mathtt{dom}(b) \rightarrowtail^{\top} u$ then

$$\mathtt{fta}_u(\alpha) \mathbin{\hat{}} \mathtt{I}(f)(\alpha, a) = \mathtt{fta}_u(\alpha a) = \mathtt{fta}_u(\beta b) = \mathtt{fta}_u(\beta) \mathbin{\hat{}} (\mathtt{fta}_u(\beta), b).$$

    Since $\mathtt{I}(f)(\alpha, a) \neq \epsilon$, it follows that $\mathtt{fta}_u(\alpha) = \mathtt{fta}_u(\beta)$ and $\mathtt{I}(f)(\alpha, a) = (\mathtt{fta}_u(\beta), b)$. However, by the assumption that the interpretation is non-conflating, the latter implies that $\mathtt{dom}(b) \neq u$. Also, $\mathtt{dom}(a) \neq u$ by reflexivity. By the induction hypothesis, we obtain from $\mathtt{fta}_u(\alpha) = \mathtt{fta}_u(\beta)$ that $\mathtt{view}_u(\alpha) = \mathtt{view}_u(\beta)$. Thus

$$\begin{aligned} \mathtt{view}_u(\alpha a) &= \mathtt{view}_u(\alpha) \circ \mathtt{obs}_u(\alpha a) \\ &= \mathtt{view}_u(\beta) \circ \mathtt{obs}_u(\beta b) \\ &= \mathtt{view}_u(\beta b). \end{aligned}$$

- If $\mathtt{dom}(a) \rightarrowtail^{\top} u$ then $\mathtt{fta}_u(\alpha') = \mathtt{fta}_u(\alpha a) = \mathtt{fta}_u(\alpha)\,\hat{}\,(\mathtt{fta}_{\mathtt{dom}(a)}(\alpha), a)$. This means that $\alpha' \neq \epsilon$. Assume $\alpha' = \beta b$.

  - If either $\mathtt{dom}(b) \not\rightarrowtail u$, or $\mathtt{dom}(b) \rightarrowtail^{g} u$ and $g \neq \top$ then we can swap the roles of αa and βb and apply the argument from the first three cases above.
  - If $\mathtt{dom}(b) \rightarrowtail^{\top} u$ then $\mathtt{fta}_u(\alpha a) = \mathtt{fta}_u(\alpha)\,\hat{}\,(\mathtt{fta}_{\mathtt{dom}(a)}(\alpha), a)$ and $\mathtt{fta}_u(\beta b) = \mathtt{fta}_u(\beta)\,\hat{}\,(\mathtt{fta}_{\mathtt{dom}(b)}(\beta), b)$. Therefore, we have $a = b$, and $\mathtt{fta}_u(\alpha) = \mathtt{fta}_u(\beta)$. By the inductive hypothesis, we have $\mathtt{view}_u(\alpha) = \mathtt{view}_u(\beta)$. If $\mathtt{dom}(a) = u$, then

$$\begin{aligned}\mathtt{view}_u(\alpha a) &= \mathtt{view}_u(\alpha)\, a\, \mathtt{obs}_u(\alpha a)\\ &= \mathtt{view}_u(\beta)\, b\, \mathtt{obs}_u(\beta b)\\ &= \mathtt{view}_u(\beta b).\end{aligned}$$

    Alternately, if $\mathtt{dom}(a) \neq u$, then

$$\begin{aligned}\mathtt{view}_u(\alpha a) &= \mathtt{view}_u(\alpha) \circ \mathtt{obs}_u(\alpha a)\\ &= \mathtt{view}_u(\beta) \circ \mathtt{obs}_u(\beta b)\\ &= \mathtt{view}_u(\beta b).\end{aligned}$$

□

We note that this result does not hold for conflating interpreted architectures. For example, consider an interpreted architecture in which $\mathtt{dom}(a) \rightarrowtail^{f} \mathtt{dom}(b)$ and $\mathtt{I}(f)(\alpha, a) = (\epsilon, b)$ and a system in which $\mathtt{obs}_u(s) = \bot$ for all $u \in D$ and states s. This system is necessarily FTA-compliant with the architecture, but we have $\mathtt{fta}_{\mathtt{dom}(b)}(a) = (\epsilon, b) = \mathtt{fta}_{\mathtt{dom}(b)}(b)$ but $\mathtt{view}_{\mathtt{dom}(b)}(a) = \bot$ and $\mathtt{view}_{\mathtt{dom}(b)}(b) = \bot b \bot$.

However, it is always possible to convert an interpretation $\mathtt{I}$ into another equivalent non-conflating interpretation $\mathtt{I}'$, by defining $\mathtt{I}'(f)(\alpha, a) = (\mathtt{I}(f)(\alpha, a), x)$ for some fixed value x that is not in A. Note that $\mathtt{I}$ and $\mathtt{I}'$ are equivalent in the sense that $\mathtt{I}(f)(\alpha, a) = \mathtt{I}(g)(\beta, b)$ if and only if $\mathtt{I}'(f)(\alpha, a) = \mathtt{I}'(g)(\beta, b)$, so the "information content" of the values of $\mathtt{I}$ and $\mathtt{I}'$ are the same. We therefore assume in what follows that interpreted architectures are non-conflating.

4.2 Example: Starlight Interactive Link

The Starlight Interactive Link [Anderson et al., 1996] provides interactive access from a high-security network to a low-security network. This allows a user on the high-security network to have windows open on her screen at differing security levels, while ensuring no high-security information goes to the low-security network.

Starlight has both hardware and software components. The hardware device is connected to both the high-security and low-security networks, and has a keyboard and mouse attached. There is a switch that can toggle between the

high-security and low-security networks; input from the mouse and keyboard are sent to the network currently selected by the switch. Starlight allows data from the low-security network to be transferred to the high-security network, but not vice versa. The Starlight software components include proxy window clients and servers to allow the windowing environment to work in the presence of the Starlight hardware.

The following diagram shows extended architecture $\mathcal{SL}$, an architecture for the Starlight Interactive Link. The architecture uses a filter function to specify what information the Starlight Interactive Link may send to the low-security network.

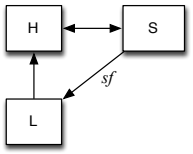

Domain H represents the high-security network (including the user's computer); domain L represents the low-security network; domain S represents the Starlight Interactive Link (including input devices), which routes keyboard and mouse events to either the high-security or low-security network. Note that there is no edge from domain L to domain S, as no information is sent directly from the low-security network to the Starlight Interactive Link. Instead, data from the low-security network (such as updates to the contents of a window) are sent to the high-security network, and thence to software components of the Starlight Interactive Link.

The edge labeled sf restricts what information is allowed to flow from the Starlight Interactive Link to low-security network L. We present an architectural specification $\mathcal{C}_{\mathcal{SL}}$ based on $\mathcal{SL}$ that expresses a constraint on interpretations of sf. An interpretation $(A, \mathtt{dom}, \mathtt{I})$ is included in $\mathcal{C}_{\mathcal{SL}}$ if the following conditions hold. Let $A_S = \{a \in A \mid \mathtt{dom}(a) = S\}$ be the set of actions belonging to domain S. We assume that there is a distinguished action $\mathtt{t} \in A_S$ that toggles which network is receiving the input events. Intuitively, L is permitted to know about the occurrence of any t action, and the occurrence of any other action in A_S (e.g., keyboard or mouse input) that happens while the low-security network is selected (i.e., after an odd number of toggle actions). We capture this by the following assumption on interpretation $\mathtt{I}$:

$$\mathtt{I}(sf)(\alpha, a) = \begin{cases} a & \text{if } a = \mathtt{t} \text{ or} \\ & \quad \#_{\mathtt{t}}(\alpha) \text{ is odd and } a \in A_S \\ \epsilon & \text{otherwise} \end{cases}$$

where $\#_{\mathtt{t}}(\alpha)$ is the number of occurrences of t in α. For example, with the α_i consisting of only H and L actions, and $a \neq \mathtt{t}$ an S action, we have

$$\mathtt{I}(sf)(\alpha_0, a) = \epsilon,$$

$\mathtt{I}(sf)(\alpha_0\, a, \mathtt{t}) = t$,
$\mathtt{I}(sf)(\alpha_0\, a\, \mathtt{t}\, \alpha_1, a) = a$, and
$\mathtt{I}(sf)(\alpha_0\, a\, \mathtt{t}\, \alpha_1\, a\, \mathtt{t}\, \alpha_2, a) = \epsilon$, etc.

It is straightforward to check for such an interpretation that it is non-conflating and that $\mathtt{I}(sf)$ is $\mathtt{fta}_S$-compatible.

4.2.1 Information security properties

The component S, corresponding to the Starlight Interactive Link, is a trusted component, and the *sf* filter is a local constraint on the component. To verify that a system satisfies specification $(\mathcal{SL}, \mathcal{C}_{\mathcal{SL}})$, we would need to verify that S appropriately filters information sent to L, and that all other communication in the system complies with the architecture, to wit, that H cannot communicate directly with L, and L cannot communicate directly with S.

If a system does satisfy specification $(\mathcal{SL}, \mathcal{C}_{\mathcal{SL}})$, then we can show that domain L never knows any H-dependent propositions. Indeed, we can show something stronger, that L never knows any proposition about H actions and S actions that occur while the high-network is selected.

To fully capture this intuition requires a little care, since a proposition such as "H did action a between the first and second $\mathtt{t}$ actions" should not be known to L, but refers both to something that should be hidden from L and to something that L is permitted to observe (the $\mathtt{t}$ actions). We handle this using an approach that is similar to the way we used G-dependent propositions above.

Let the *canonical form* of a sequence $\alpha \in A^*$ be the (unique) representation $\alpha = \alpha_0 \mathtt{t} \alpha_1 \mathtt{t} \alpha_2 \ldots \mathtt{t} \alpha_n$ where each α_i contains no $\mathtt{t}$ actions. Define $togL : A^* \to A^*$ be the function such that if $\alpha = \alpha_0 \mathtt{t} \alpha_1 \mathtt{t} \alpha_2 \ldots \mathtt{t} \alpha_n$ is the canonical form, then

$$togL(\alpha) = (\alpha_0 \upharpoonright L)\, \mathtt{t}\, (\alpha_1 \upharpoonright LS)\, \mathtt{t}\, (\alpha_2 \upharpoonright L)\, \mathtt{t} \ldots \mathtt{t}\, (\alpha_n \upharpoonright H_n),$$

where $H_n = L$ if n is even and $H_n = LS$ if n is odd. Intuitively, this is the subsequence consisting of the events that L is permitted to observe. All other events in α are H events, and S events that occurred while the system was toggled to High. We say that a proposition X is *toggle-High dependent at* α if there exists a sequence $\beta \in A^*$ such that $togL(\alpha) = togL(\beta)$ but $\alpha \in X$ iff $\beta \notin X$. Intuitively, this says that changing iα by adding or deleting H events and S events that L is not permitted to know about, we can change the value of the proposition X.[3]

Theorem 5 *Suppose system M is FTA-compliant with architectural specification $(\mathcal{SL}, \mathcal{C}_{\mathcal{SL}})$. If $\pi(p)$ is a toggle-H dependent proposition at α, then $M, \pi, \alpha \vDash \neg K_L p$.*

[3]This definition of toggle-High-dependent differs from that used in the earlier version of this work published at LAW12. The present formulation simplifies and strengthens the corresponding result, Theorem 5.

Proof: We first show that $\mathtt{fta}_L(\alpha) = \mathtt{fta}_L(togL(\alpha))$ for all $\alpha \in A^*$. The proof is by induction on α. The base case of $\alpha = \epsilon$ is trivial. Consider the sequence αa and assume that $\mathtt{fta}_L(\alpha) = \mathtt{fta}_L(togL(\alpha))$. Let the canonical form of α be $\alpha = \alpha_0 \mathtt{t} \alpha_1 \mathtt{t} \alpha_2 \ldots \mathtt{t} \alpha_n$

If $\mathtt{dom}(a) = L$ then $togL(\alpha a) = togL(\alpha)a$ and

$$\begin{aligned}
\mathtt{fta}_L(\alpha a) &= \mathtt{fta}_L(\alpha)\,\hat{}\,(\mathtt{fta}_L(\alpha), a) \\
&= \mathtt{fta}_L(togL(\alpha))\,\hat{}\,(\mathtt{fta}_L togL((\alpha)), a) && \text{by induction} \\
&= \mathtt{fta}_L(togL(\alpha)a) \\
&= \mathtt{fta}_L(togL(\alpha a))
\end{aligned}$$

as required. If $\mathtt{dom}(a) = S$ and $\#_{\mathtt{t}}(\alpha)$ is odd then $togL(\alpha a) = togL(\alpha)a$ and

$$\begin{aligned}
\mathtt{fta}_L(\alpha a) &= \mathtt{fta}_L(\alpha)\,\hat{}\,\mathtt{I}(sf)(\alpha, a) \\
&= \mathtt{fta}_L(togL(\alpha))\,\hat{}\,\mathtt{I}(sf)(\alpha, a) && \text{by induction} \\
&= \mathtt{fta}_L(togL(\alpha))\,\hat{}\,\mathtt{I}(sf)(togL(\alpha), a) && \text{by observation above} \\
&= \mathtt{fta}_L(togL(\alpha)a) \\
&= \mathtt{fta}_L(togL(\alpha a))
\end{aligned}$$

as required. If $\mathtt{dom}(a) = S$ and $\#_{\mathtt{t}}(\alpha)$ is even then $togL(\alpha a) = togL(\alpha)$ and $\mathtt{I}(sf)(\alpha, a) = \epsilon$, so

$$\begin{aligned}
\mathtt{fta}_L(\alpha a) &= \mathtt{fta}_L(\alpha)\,\hat{}\,\mathtt{I}(sf)(\alpha, a) \\
&= \mathtt{fta}_L(\alpha) \\
&= \mathtt{fta}_L(togL(\alpha)) && \text{by induction} \\
&= \mathtt{fta}_L(togL(\alpha a))
\end{aligned}$$

as required. Finally, if $\mathtt{dom}(a) = H$ then $togL(\alpha a) = togL(\alpha)$, and since $H \not\rightarrow L$, we have

$$\begin{aligned}
\mathtt{fta}_L(\alpha a) &= \mathtt{fta}_L(\alpha) \\
&= \mathtt{fta}_L(togL(\alpha)) && \text{by induction} \\
&= \mathtt{fta}_L(togL(\alpha a))
\end{aligned}$$

as required. This completes the proof that $\mathtt{fta}_L(\alpha) = \mathtt{fta}_L(togL(\alpha))$.

We now prove the theorem. Suppose system M is FTA-compliant with architectural specification $(\mathcal{SL}, \mathcal{C}_{\mathcal{SL}})$ and let $\pi(p)$ be a toggle-H dependent proposition at α. Them there exists $\beta \in A^*$ such that $togL(\alpha) = togL(\beta)$ and $\alpha \in \pi(p)$ iff $\beta \notin \pi(p)$. By what was shown above, we have $\mathtt{fta}_L(\alpha) = \mathtt{fta}_L(\beta)$.

Since M is FTA-compliant, it follows that $\mathtt{view}_L(\alpha) = \mathtt{view}_L(\beta)$, and we have $M, \pi, \alpha \models \neg K_L p$. □

Note that we would not be able show this security property in an architecture without filter functions, since the domain S can communicate with both H and L. It is the filter function that allows us to show that S's communication with L reveals nothing about H actions, and only limited information about S actions.

4.3 Example: Downgrader

Many systems downgrade information, releasing confidential information to untrusted entities. Because it is a sensitive operation, downgrading is typically restricted to certain trusted components, called downgraders.

The following diagram shows extended architecture $\mathcal{DG}$, containing low-security domain L, high-security data store H for two domains of high-security users C and P, and downgrader D.

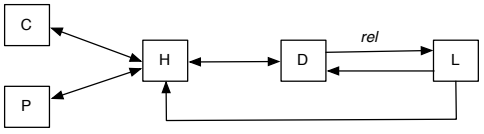

This architecture represents a system where some information is allowed to flow from the high-security domain H to low-security domain L, but only through the downgrader D. In particular, we assume that domain L is permitted to know about the actions of P (via downgrader D), but should never know anything about the actions of C. This may be an appropriate model for a government agency where some sensitive information may be released to the public, but certain sensitive data (e.g., the identity and activity of covert employees, represented by domain C) should never be released.

The filter function for *rel* restricts the information that may be released from D to L. It is a local constraint on the trusted component D. The filter function should ensure that nothing is revealed about the actions of C. We define an architectural specification $\mathcal{C}_{\mathcal{DG}}$ to restrict our attention to architectural interpretations with suitable filter functions.

We first define an interpretation of the function *rel* that states that the information transmitted across this edge is the maximum information that D would have if the domain C were completely cut off from the other domains. Let $\rightarrowtail'$ be the information flow policy such that $(u, v, f) \in \rightarrowtail'$ if and only if $f = \top$ and either $u = v = C$ or there exists an edge $(u, v, g) \in \rightarrowtail$ and $\{u, v\} \subseteq \{P, H, D, L\}$. This is the policy depicted in the following diagram, where, as usual, we omit reflexive edges:

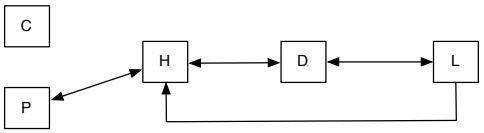

Let $(A, \mathtt{dom}, \mathtt{I}) \in \mathcal{C}_{\mathcal{DG}}$ if and only if $(A, \mathtt{dom}, \mathtt{I})$ is an interpretation of $\mathcal{DG}$ such that $\mathtt{I}(rel)(\alpha, a) = (\mathtt{fta}^{\rightarrowtail'}_{D}(\alpha), a)$ for all $\alpha \in A^*$ and $a \in A$ with $\mathtt{dom}(a) = D$. Thus, architectural interpretations in $\mathcal{C}_{\mathcal{DG}}$ ensure that $\mathtt{fta}^{\rightarrowtail'}_{D}$ is an upper-bound on information that may be released from D to L. It is straightforward to check that these architectural specifications are non-conflating. The required compatibility constraint is also satisfied, as shown in the following result.

Proposition 1 *If $(A, \mathtt{dom}, \mathtt{I}) \in \mathcal{C}_{\mathcal{DG}}$ then $\mathtt{I}(rel)$ is $\mathtt{fta}^{\rightarrowtail}_{D}$-compatible.*

Proof: For $u \in \{P, H, D, L\}$, define the functions F_u inductively, by $F_u(\epsilon) = \epsilon$ and

$$F_u(\sigma \hat{} (\delta, a)) = \begin{cases} F_u(\sigma) & \text{if } dom(a) = C \\ F_u(\sigma) \hat{} (F_{\mathtt{dom}(a)}(\delta), a) & \text{if } dom(a) \neq C \text{ and } \mathtt{dom}(a) \neq D \text{ and } u \neq L \\ F_u(\sigma) \hat{} (\delta, a) & \text{if } \mathtt{dom}(a) = D \text{ (so } dom(a) \neq C) \text{ and } u = L \end{cases}$$

We claim that for all $\alpha \in A^*$ and $u \in \{P, H, D, L\}$, we have $F_u(\mathtt{fta}^{\rightarrowtail}_{u}(\alpha)) = \mathtt{fta}^{\rightarrowtail'}_{u}(\alpha)$. It is then immediate that $\mathtt{fta}^{\rightarrowtail}_{D}(\alpha) = \mathtt{fta}^{\rightarrowtail}_{D}(\beta)$ implies $\mathtt{fta}^{\rightarrowtail'}_{D}(\alpha) = \mathtt{fta}^{\rightarrowtail'}_{D}(\beta)$, and it easily follows that $\mathtt{I}(rel)$ is $\mathtt{fta}^{\rightarrowtail}_{D}$-compatible.

To prove the claim we proceed by induction on α. the base case of $\alpha = \epsilon$ is trivial. Assuming the claim holds for α, let $u \in \{P, H, D, L\}$ and consider the sequence αa. There are several cases:

- Case 1: $\mathtt{dom}(a) \rightarrowtail^{\top} u$ and $\mathtt{dom}(a) \neq C$. By definition of $\rightarrowtail'$, we have $\mathtt{dom}(a) \rightarrowtail'^{\top} u$. Thus,

$$\begin{aligned} F_u(\mathtt{fta}_u(\alpha a)) &= F_u(\mathtt{fta}_u(\alpha) \hat{} (\mathtt{fta}_{\mathtt{dom}(a)}(\alpha), a)) \\ &= F_u(\mathtt{fta}_u(\alpha)) \hat{} (F_{\mathtt{dom}(a)}(\mathtt{fta}_{\mathtt{dom}(a)}(\alpha)), a) \\ &= \mathtt{fta}^{\rightarrowtail'}_{u}(\alpha) \hat{} (\mathtt{fta}^{\rightarrowtail'}_{\mathtt{dom}(a)}(\alpha), a) && \text{by induction} \\ &= \mathtt{fta}^{\rightarrowtail'}_{u}(\alpha a). \end{aligned}$$

- Case 2: $\mathtt{dom}(a) \rightarrowtail^{\top} u$ and $\mathtt{dom}(a) = C$. In this case, $\mathtt{dom}(a) \not\rightarrowtail' u$. Thus,

$$\begin{aligned} F_u(\mathtt{fta}_u(\alpha a)) &= F_u(\mathtt{fta}_u(\alpha) \hat{} (\mathtt{fta}_{\mathtt{dom}(a)}(\alpha), a)) \\ &= F_u(\mathtt{fta}_u(\alpha)) \\ &= \mathtt{fta}^{\rightarrowtail'}_{u}(\alpha) && \text{by induction} \\ &= \mathtt{fta}^{\rightarrowtail'}_{u}(\alpha a). \end{aligned}$$

- Case 3: $\mathtt{dom}(a) \overset{rel}{\rightarrowtail} u$, so $\mathtt{dom}(a) = D$ and $u = L$. By definition of $\rightarrowtail'$, we have $\mathtt{dom}(a) \overset{\top}{\rightarrowtail}' u$. Thus,

$$\begin{aligned} F_u(\mathtt{fta}_u(\alpha a)) &= F_u(\mathtt{fta}_u(\alpha) \hat{} \mathtt{I}(rel)(\alpha, a) \\ &= F_u(\mathtt{fta}_u(\alpha) \hat{} (\mathtt{fta}^{\rightarrowtail'}_{\mathtt{dom}(a)}(\alpha), a)) \\ &= F_u(\mathtt{fta}_u(\alpha)) \hat{} (\mathtt{fta}^{\rightarrowtail'}_{\mathtt{dom}(a)}(\alpha), a) && \text{by def. } F_u \\ &= \mathtt{fta}^{\rightarrowtail'}_u(\alpha) \hat{} (\mathtt{fta}^{\rightarrowtail'}_{\mathtt{dom}(a)}(\alpha), a) && \text{by induction} \\ &= \mathtt{fta}^{\rightarrowtail'}_u(\alpha a). \end{aligned}$$

- Case 4: $\mathtt{dom}(a) \not\rightarrowtail u$. In this case also $\mathtt{dom}(a) \not\rightarrowtail' u$. Thus,

$$\begin{aligned} F_u(\mathtt{fta}_u(\alpha a)) &= F_u(\mathtt{fta}_u(\alpha)) \\ &= \mathtt{fta}^{\rightarrowtail'}_u(\alpha) && \text{by induction} \\ &= \mathtt{fta}^{\rightarrowtail'}_u(\alpha a). \end{aligned}$$

  Thus, the claim holds in all cases.

□

The architectural specification imposes restrictions on what information L knows and when. Domain L can only learn about other domains via downgrader D. The restriction on filter functions for *rel* ensures that L cannot know anything about the actions of C, since $C \not\rightarrowtail' H$, and so the function $\mathtt{fta}^{\rightarrowtail'}_D$ does not contain any information about the actions of C. However, information about the actions of H, P and D may be released to L. Also, L cannot learn about actions of other domains that occurred after the last D action. The following two theorems express these restrictions.

Theorem 6 *If M is FTA-compliant with architectural specification $(\mathcal{DG}, \mathcal{C}_{\mathcal{DG}})$ and $\pi(p)$ depends on G actions at α, then $M, \pi, \alpha \vDash \neg K_L p$.*

Proof:

Let $G = \overline{C} = \{P, H, D, L\}$. We first claim that $\mathtt{fta}^{\rightarrowtail'}_u(\gamma \restriction G) = \mathtt{fta}^{\rightarrowtail'}_u(\gamma)$ for any $\gamma \in A^*$ and $u \in G$. We proceed by induction on $|\gamma|$. The base case is trivial. Suppose that $\mathtt{fta}^{\rightarrowtail'}_u(\gamma \restriction G) = \mathtt{fta}^{\rightarrowtail'}_u(\gamma)$ for $u \in G$ and consider γa. We consider several cases.

- If $\mathtt{dom}(a) \overset{f}{\rightarrowtail}' u$ then $f = \top$ and $\mathtt{dom}(a) \in G$, by the definition of $\rightarrowtail'$. Thus

$$\begin{aligned} \mathtt{fta}^{\rightarrowtail'}_u((\gamma a) \restriction G) &= \mathtt{fta}^{\rightarrowtail'}_u((\gamma \restriction G)a) \\ &= \mathtt{fta}^{\rightarrowtail'}_D(\gamma \restriction G) \hat{} (\mathtt{fta}_{\mathtt{dom}(a)}(\gamma \restriction G), a) && \text{by defn } \mathtt{fta}^{\rightarrowtail'}_D \\ &= \mathtt{fta}^{\rightarrowtail'}_D(\gamma) \hat{} (\mathtt{fta}_{\mathtt{dom}(a)}(\gamma), a) && \text{by IH} \\ &= \mathtt{fta}^{\rightarrowtail'}_D(\gamma a) && \text{by defn } \mathtt{fta}^{\rightarrowtail'}_D \end{aligned}$$

- If $\mathtt{dom}(a) \not\rightarrowtail' u$ and $\mathtt{dom}(a) \in G$ then

$$\begin{aligned}
\mathtt{fta}_u^{\rightarrowtail'}((\gamma a) \restriction G) &= \mathtt{fta}_u^{\rightarrowtail'}((\gamma \restriction G)a) \\
&= \mathtt{fta}_u^{\rightarrowtail'}(\gamma \restriction G) && \text{by defn } \mathtt{fta}_u^{\rightarrowtail'} \\
&= \mathtt{fta}_u^{\rightarrowtail'}(\gamma) && \text{by IH} \\
&= \mathtt{fta}_u^{\rightarrowtail'}(\gamma a) && \text{by defn } \mathtt{fta}_u^{\rightarrowtail'}
\end{aligned}$$

- If $\mathtt{dom}(a) \not\rightarrowtail' u$ and $\mathtt{dom}(a) \notin G$ then

$$\begin{aligned}
\mathtt{fta}_u^{\rightarrowtail'}((\gamma a) \restriction G) &= \mathtt{fta}_u^{\rightarrowtail'}(\gamma \restriction G) && \text{as } \mathtt{dom}(a) \notin G \\
&= \mathtt{fta}_u^{\rightarrowtail'}(\gamma) && \text{by IH} \\
&= \mathtt{fta}_u^{\rightarrowtail'}(\gamma a) && \text{by defn } \mathtt{fta}_u^{\rightarrowtail'}
\end{aligned}$$

This completes the proof of the claim. In particular, we have $\mathtt{fta}_D^{\rightarrowtail'}(\gamma \restriction G) = \mathtt{fta}_D^{\rightarrowtail'}(\gamma)$, for all γ. Note that this implies, for all γ and $a \in A$ with $\mathtt{dom}(a) = D$, that $\mathtt{I}(rel)(\gamma \restriction G, a) = (\mathtt{fta}_D^{\rightarrowtail'}(\gamma \restriction G), a) = (\mathtt{fta}_D^{\rightarrowtail'}(\gamma), a) = \mathtt{I}(rel)(\gamma, a)$.

We now show that $\mathtt{fta}_L(\gamma) = \mathtt{fta}_L(\gamma \restriction G)$. We proceed by induction on $|\gamma|$. The base case is trivial. Consider γa, and assume that $\mathtt{fta}_L(\gamma) = \mathtt{fta}_L(\gamma \restriction G)$. We consider several cases.

- If $\mathtt{dom}(a) \in \{C, P, H\}$ then $\mathtt{dom}(a) \not\rightarrowtail L$, so $\mathtt{fta}_L(\gamma a) = \mathtt{fta}_L(\gamma)$ by definition of $\mathtt{fta}_L$. We have that $(\gamma a) \restriction G$ is either $\gamma \restriction G$ or $(\gamma \restriction G)a$, but in either case, $\mathtt{fta}_L((\gamma a) \restriction G) = \mathtt{fta}_L(\gamma \restriction G)$ by the definition of $\mathtt{fta}_L$. By the inductive hypothesis, we have $\mathtt{fta}_L(\gamma a) = \mathtt{fta}_L((\gamma a) \restriction G)$ as required.

- If $\mathtt{dom}(a) = L$ then $\gamma a \restriction G = (\gamma \restriction G)a$, so

$$\begin{aligned}
\mathtt{fta}_L(\gamma a) &= \mathtt{fta}_L(\gamma)\,\hat{}\,(\mathtt{fta}_L(\gamma), a) \\
&= \mathtt{fta}_L(\gamma \restriction G)\,\hat{}\,(\mathtt{fta}_L(\gamma \restriction G), a) && \text{by the IH} \\
&= \mathtt{fta}_L((\gamma \restriction G)a) \\
&= \mathtt{fta}_L((\gamma a) \restriction G)
\end{aligned}$$

- If $\mathtt{dom}(a) = D$ then

$$\begin{aligned}
\mathtt{fta}_L(\gamma a) &= \mathtt{fta}_L(\gamma)\,\hat{}\,\mathtt{I}(rel)(\gamma, a) \\
&= \mathtt{fta}_L(\gamma)\,\hat{}\,\mathtt{I}(rel)(\gamma \restriction G, a) && \text{as } \mathtt{I}(rel)(\gamma \restriction G, a) = \mathtt{I}(rel)(\gamma, a) \\
&= \mathtt{fta}_L(\gamma \restriction G)\,\hat{}\,\mathtt{I}(rel)(\gamma \restriction G, a) && \text{by the IH} \\
&= \mathtt{fta}_L((\gamma \restriction G)a) \\
&= \mathtt{fta}_L((\gamma a) \restriction G) && \text{since } \mathtt{dom}(a) = D \in G
\end{aligned}$$

This completes the proof that $\mathtt{fta}_L(\gamma) = \mathtt{fta}_L(\gamma \restriction G)$ for all γ.

Since $\pi(p)$ depends on C actions at α, there exists $\beta \in A^*$ such that $\alpha \restriction G = \beta \restriction G$ but $\alpha \in \pi(p)$ iff $\beta \notin \pi(p)$. By the claim just proved, we have

$\mathtt{fta}_L(\alpha) = \mathtt{fta}_L(\alpha \restriction G) = \mathtt{fta}_L(\beta \restriction G) = \mathtt{fta}_L(\beta)$, from which $\alpha \approx_L \beta$ follows by compliance. It follows that $M\,\pi, \alpha \models \neg K_L(p)$. □

Actions of the downgrader are the only means by which L acquires knowledge about other domains' actions. Thus, L knows nothing about any action of other domains that occur after the last D action. To state this formally, we say that proposition X is *about C, P, and H activity after the last D action* if there exists a set $Y \subseteq A^*$ such that for all $\alpha \in A^*$, $\alpha \in X \iff F(\alpha) \in Y$, where F is defined recursively by $F(\epsilon) = \epsilon$ and

$$F(\beta a) = \begin{cases} \epsilon & \text{if } \mathtt{dom}(a) = D \\ F(\beta)a & \text{if } \mathtt{dom}(a) \in \{C, P, H\} \\ F(\beta) & \text{otherwise} \end{cases}$$

Intuitively, F returns the C, P, and H actions that occur after the last D action, and so proposition X is about C, P, and H activity after the last D action if membership in X depends only on the C, P, and H actions that occur after the last D action.

Theorem 7 *If M is FTA-compliant with architectural specification $(\mathcal{DG}, \mathcal{C}_{\mathcal{DG}})$ and $\pi(p)$ is a non-trivial proposition about C, P, and H activity after the last D action, then $M, \pi \vDash \neg K_L p$.*

Proof: Let $\alpha \in A^*$. Since $\pi(p)$ is non-trivial, there is an α' such that $\alpha \in \pi(p) \iff \alpha' \notin \pi(p)$. Let α_1 be the longest suffix of α such that α_1 does not contain any D actions. Then $\alpha = \alpha_0\alpha_1$, where either $\alpha_0 = \epsilon$ or the last action of α_0 is a D action. Let $\beta = \alpha_0(\alpha_1 \restriction \{L\})F(\alpha')$. Note that $F(\beta) = F(\alpha')$, and so $\alpha \in \pi(p) \iff \beta \notin \pi(p)$. We will show that $\mathtt{fta}_L(\alpha) = \mathtt{fta}_L(\beta)$.

First, note that $\mathtt{fta}_L(\beta) = \mathtt{fta}_L(\alpha_0(\alpha_1 \restriction \{L\})F(\alpha')) = \mathtt{fta}_L(\alpha_0(\alpha_1 \restriction \{L\}))$ since $F(\alpha')$ contains only actions a such that $\mathtt{dom}(a) \not\rightarrow L$. Note also that $\mathtt{fta}_L(\alpha) = \mathtt{fta}_L(\alpha_0\alpha_1) = \mathtt{fta}_L(\alpha_0(\alpha_1 \restriction \{L\}))$ since α_1 contains only L actions and actions a such that $\mathtt{dom}(a) \not\rightarrow L$.

Therefore, $\mathtt{fta}_L(\beta) = \mathtt{fta}_L(\alpha)$ as required. Since M is FTA-compliant, by Lemma 4 we have $\mathtt{view}_L(\alpha) = \mathtt{view}_L(\beta)$, and so $\alpha \approx_L \beta$, which immediately implies the result. □

4.4 Example: Electronic election

The following diagram shows architecture $\mathcal{EE}$, an electronic election for n voters, coordinated by an election authority *ElecAuth*.

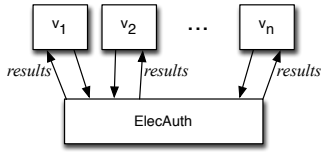

There is one domain v_i for each voter, and a single domain *ElecAuth* for the election authority. The architecture permits arbitrary information to flow from a voter to the election authority, but uses filter function *results* to restrict what information may flow from the election authority to each voter.

In many elections, the behavior of individual voters is confidential information: it should not be known how voters voted. (Elections have several security requirements. For example, the final result should be correctly computed from the votes—an integrity requirement. We focus here on the confidentiality requirement for voters.)

By specifying an additional local constraint on the election authority, we can show that this architecture enforces anonymity on the identity of the voters, thus satisfying the confidentiality requirement about the behavior of voters. The election authority's compliance with this local constraint might be assured by means of a careful verification of its implementation, or carefully designed cryptographic protocols (which may remove some or all of the trust required to be placed in the election authority).

We first assume that voters are homogenous in that they have the same set of actions. Let $(\mathcal{EE}, \mathcal{I})$ be an interpreted architecture where $\mathcal{I} = (A, \texttt{dom}, \texttt{I})$. We say that architecture interpretation $\mathcal{I}$ is *voter homogenous* if for any voter v, the set of possible actions $A_v = \{a \in A \mid \texttt{dom}(a) = v\}$ equals $\{a^v \mid a \in A_V\}$, where A_V is the set of action types available to voters. Intuitively, a^v represent the action of type a when performed by voter v.

We define a voter permutation P as a permutation over the set of voters V. Since all voters have the same set of actions in a voter-homogenous interpretation, we can apply a permutation P to sequence $\alpha \in A^*$, written $P(\alpha)$, as follows:

$$\begin{aligned} P(\epsilon) &= \epsilon \\ P(\alpha a) &= P(\alpha)a && \text{if } \texttt{dom}(a) \notin V \\ P(\alpha a^v) &= P(\alpha)a^{P(v)} && \text{if } a \in A_V \ . \end{aligned}$$

We apply permutation P to proposition $X \subseteq A^*$ by applying P to each sequence $\alpha \in X$, i.e.,

$$P(X) = \{P(\alpha) \mid \alpha \in X\} \ .$$

Using voter permutations, we can now state the local constraint that election results do not depend on (and thus do not reveal) the identity of any voter.

A1. Election results have the following *identity-oblivious* property: Given voter-homogenous interpreted architecture $(\mathcal{EE}, \mathcal{I})$ where $\mathcal{I} = (A, \texttt{dom}, \texttt{I})$, for all voter permutations P, sequences $\alpha \in A^*$, and actions $a \in A$ with $\texttt{dom}(a) = \mathit{ElecAuth}$, we have $\texttt{I}(\mathit{results})(\alpha, a) = \texttt{I}(\mathit{results})(P(\alpha), a)$.

There are several possible interpretations of *results* that satisfy this constraint, such as a function that returns each candidate and her vote tally, or a function that returns the total number of votes submitted. However, a function

that returns a ballot and the identity of the voter that submitted it doesn't satisfy constraint A1.

We define architecture specification $\mathcal{C}_{\mathcal{EE}}$ such that $\mathcal{I} \in \mathcal{C}_{\mathcal{EE}}$ if and only if $\mathcal{I}$ is a voter-homogenous interpretation of $\mathcal{EE}$ that satisfies constraint A1.

Given the local constraint A1, we can show that if voter v believes that some proposition X may be satisfied, then v also believes that $P(X)$ may be satisfied, for any voter permutation P with $P(v) = v$. For example, if Alice considers it possible that Bob voted for Obama and Charlie voted for Romney, then Alice considers it possible that Charlie voted for Obama and Bob voted for Romney.

Before stating and proving this result, we first prove a useful lemma: a voter's view of an election does not depend on the identity of other voters. That is, for any permutation P such that $P(v) = v$, the sequences α and $P(\alpha)$ appear the same to voter v.

Lemma 5 *Let system M be FTA-compliant with $(\mathcal{EE}, \mathcal{C}_{\mathcal{EE}})$. For all voters v, voter permutations P such that $P(v) = v$, and sequences α, $\alpha \approx_v P(\alpha)$.*

Proof: Since M is FTA-compliant with $(\mathcal{EE}, \mathcal{C}_{\mathcal{EE}})$, there is an $\mathcal{I} = (A, \texttt{dom}, \texttt{I}) \in \mathcal{C}_{\mathcal{EE}}$ such that M is FTA-compliant with interpreted architecture $(\mathcal{EE}, \mathcal{I})$.

We show, by induction on the length of α, that $\texttt{fta}_v(\alpha) = \texttt{fta}_v(P(\alpha))$. The base case is trivial. Consider sequence αa, and assume that $\texttt{fta}_v(\alpha) = \texttt{fta}_v(P(\alpha))$. We show that $\texttt{fta}_v(\alpha a) = \texttt{fta}_v(P(\alpha a))$.

If $\texttt{dom}(a) = v$ then

$$\begin{aligned}
\texttt{fta}_v(P(\alpha a)) &= \texttt{fta}_v(P(\alpha)a) && \text{as } P(a) = a \\
&= \texttt{fta}_v(P(\alpha))\,\hat{}\,(\texttt{fta}_v(P(\alpha)), a) && \text{by defn } \texttt{fta}_v \\
&= \texttt{fta}_v(\alpha)\,\hat{}\,(\texttt{fta}_v(\alpha), a) && \text{by IH} \\
&= \texttt{fta}_v(\alpha a) && \text{by defn } \texttt{fta}_v
\end{aligned}$$

If $\texttt{dom}(a) \in V \setminus \{v\}$ then $\texttt{dom}(a) \not\rightarrow v$, so

$$\begin{aligned}
\texttt{fta}_v(P(\alpha a)) &= \texttt{fta}_v(P(\alpha)b) && \text{for some } b \in A,\ \texttt{dom}(b) \in V \setminus \{v\} \\
&= \texttt{fta}_v(P(\alpha)) && \text{by defn } \texttt{fta}_v \\
&= \texttt{fta}_v(\alpha) && \text{by IH} \\
&= \texttt{fta}_v(\alpha a) && \text{by defn } \texttt{fta}_v
\end{aligned}$$

If $\texttt{dom}(a) = \mathit{ElecAuth}$ then

$$\begin{aligned}
\texttt{fta}_v(P(\alpha a)) &= \texttt{fta}_v(P(\alpha)a) && \text{as } P(a) = a \\
&= \texttt{fta}_v(P(\alpha))\,\hat{}\,\texttt{I}(\mathit{results})(P(\alpha), a) && \text{by defn } \texttt{fta}_v \\
&= \texttt{fta}_v(\alpha)\,\hat{}\,\texttt{I}(\mathit{results})(P(\alpha), a)) && \text{by IH} \\
&= \texttt{fta}_v(\alpha)\,\hat{}\,\texttt{I}(\mathit{results})(\alpha, a) && \text{by A1} \\
&= \texttt{fta}_v(\alpha a) && \text{by defn } \texttt{fta}_v
\end{aligned}$$

Thus, if all cases, $\texttt{fta}_v(\alpha) = \texttt{fta}_v(P(\alpha))$. Therefore, by Lemma 4, $\alpha \approx_v P(\alpha)$. □

Theorem 8 *Let system M be FTA-compliant with architectural specification $(\mathcal{EE}, \mathcal{C}_{\mathcal{EE}})$. Let v be a voter. For all voter permutations P such that $P(v) = v$, if $\pi(q) = P(\pi(p))$, then $M, \pi, \alpha \vDash \neg K_v \neg p \Rightarrow \neg K_v \neg q$ for all $\alpha \in A^*$.*

Proof: Suppose $M, \pi, \alpha \vDash \neg K_v \neg p$. Then there is some sequence β such that $\mathtt{view}_v(\alpha) = \mathtt{view}_v(\beta)$ and $\beta \in \pi(p)$. By Lemma 5, $\mathtt{view}_v(\beta) = \mathtt{view}_v(P(\beta))$, so $\mathtt{view}_v(\alpha) = \mathtt{view}_v(P(\beta))$. Since $P(\beta) \in P(\pi(p))$, we have $M, \pi, P(\beta) \vDash q$, so $M, \pi, \alpha \vDash \neg K_v \neg q$ as required. □

Note that Theorem 8 does not imply that voter v learns nothing about other voters. For example, if the election results reveal that all voters voted for Obama, then voter v knows how every other voter voted. Also, if there are only 2 voters, v and w, then the results may reveal to v exactly how w voted. However, Theorem 8 provides anonymity: given the results, voter v cannot distinguish the behavior of other voters.

5 Architectural refinement

During the process of system development, architectural designs are often refined by decomposing components into sub-components. We investigate the preservation of information security properties with respect to architectural refinement, including refinement of extended architectures.

Van der Meyden [2012] defines architectural refinement for information-flow architectures as follows: architecture $\mathcal{A}_1 = (D_1, \rightarrowtail_1)$ is a refinement of architecture $\mathcal{A}_2 = (D_2, \rightarrowtail_2)$ if there is a function $r : D_1 \to D_2$ such that r is onto D_2 and for all $u, v \in D_1$, if $u \rightarrowtail_1 v$ then $r(u) \rightarrowtail_2 r(v)$. We write $\mathcal{A}_1 \leq_r \mathcal{A}_2$ if $\mathcal{A}_1$ refines $\mathcal{A}_2$ via refinement function r.

Intuitively, if $\mathcal{A}_1$ refines $\mathcal{A}_2$, then $\mathcal{A}_1$ provides more detail than $\mathcal{A}_2$, as it is a finer-grain specification of security domains and the information flows between them. Refinement function r indicates how domains in D_2 are decomposed into subdomains in D_1: for all $u \in D_1$, u is a subdomain of $r(u) \in D_2$. (Also, we say that $r(u)$ is the superdomain of domain u.) The refinement function r ensures that information flows between subdomains in $\mathcal{A}_1$ are in accordance with information flows between domains in $\mathcal{A}_2$.

Example: The Hinke-Schaefer architecture $\mathcal{HS}$ refines architecture $\mathcal{HL}$ via refinement function r that maps H_{user}, H_{DBMS}, and H_F to H, and maps L_{user}, L_{DBMS}, and L_F to L. All information flows within $\mathcal{HS}$ are permitted by $\mathcal{HL}$, including the flow from L_F to H_{DBMS}, which is permitted as $L \rightarrowtail H$. Figure 2 show the architectural refinement graphically. □

It is shown by van der Meyden [2012] that TA-compliance is preserved by refinement. That is, if machine $M = \langle S, s_0, A, D_1, \mathtt{step}, \mathtt{obs}, \mathtt{dom} \rangle$ is TA-compliant with architecture $\mathcal{A}_1$, and $\mathcal{A}_1 \leq_r \mathcal{A}_2$, then $r(M)$ is TA-compliant with $\mathcal{A}_2$. (Recall that $r(M)$ is the system M with the domains abstracted via the function r, see Section 3.)

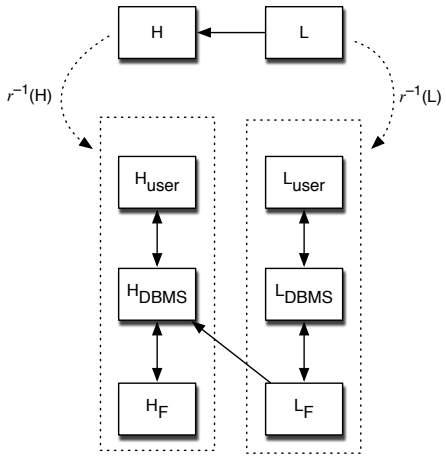

Figure 2: The Hinke-Schaefer architecture $\mathcal{HS}$ refines architecture $\mathcal{HL}$. The dashed lines and rectangles indicate how the refinement function maps domains in $\mathcal{HS}$ to the H and L domains of $\mathcal{HL}$.

5.1 Architectural refinement for interpreted extended architectures

We generalize architectural refinement to extended architectures. Since the meaning of an extended architecture relies upon the interpretation of filter functions provided in an architectural interpretation, the generalization is stated semantically, at the level of interpreted extended architectures. However, we later develop some more syntactic sufficient conditions, which may be more convenient for giving proofs that this semantic refinement relation holds.

Let $\mathcal{AI}_1 = (\mathcal{A}_1, \mathcal{I}_1)$ and $\mathcal{AI}_2 = (\mathcal{A}_2, \mathcal{I}_2)$ be interpreted extended architectures, where $\mathcal{A}_i = (D_i, \rightarrowtail_i)$, $\mathcal{I}_i = (A, \mathtt{dom}_i, \mathtt{I}_i)$ for $i = 1, 2$. (We assume that the interpretations have the same set of actions A—we do not attempt to deal with action refinement.)

Definition 5 *We say that a mapping $r : D_1 \to D_2$ is a* semantic refinement mapping *from $\mathcal{AI}_1$ to $\mathcal{AI}_2$, and write $\mathcal{AI}_1 \preceq_r \mathcal{AI}_2$, if r is onto D_2, we have $\mathtt{dom}_2 = r \circ \mathtt{dom}_1$, and for all $u \in D_1$ and sequences $\alpha, \beta \in A^*$, if $\mathtt{fta}_{r(u)}^{\mathcal{AI}_2}(\alpha) = \mathtt{fta}_{r(u)}^{\mathcal{AI}_2}(\beta)$, then $\mathtt{fta}_u^{\mathcal{AI}_1}(\alpha) = \mathtt{fta}_u^{\mathcal{AI}_1}(\beta)$.*

This definition directly implies preservation of FTA-compliance:

Theorem 9 *If $\mathcal{AI}_1 \preceq_r \mathcal{AI}_2$, then for all machines M, if M is FTA-compliant with $\mathcal{AI}_1$, then $r(M)$ is FTA-compliant with $\mathcal{AI}_2$.*

Proof: Suppose $\mathcal{AI}_1 \preceq_r \mathcal{AI}_2$. Then for all $u_2 \in D_2$, if we have $\mathtt{fta}_{u_2}(\alpha) = \mathtt{fta}_{u_2}(\alpha')$, then $\mathtt{fta}_{v_1}(\alpha) = \mathtt{fta}_{v_1}(\alpha')$ for all $v_1 \in r^{-1}(u_2) \subseteq D_1$. Since M is FTA-compliant with $\mathcal{AI}_1$, this will imply that $\mathtt{obs}_{v_1}(\alpha) = \mathtt{obs}_{v_1}(\alpha')$ for all

$v_1 \in r^{-1}(u_2)$, which by definition of $\mathtt{obs}'$ will imply that $\mathtt{obs}'_{u_2}(\alpha) = \mathtt{obs}'_{u_2}(\alpha')$, thus showing that $r(M)$ is FTA-compliant with respect to $\mathcal{AI}_2$. □

5.1.1 Sufficient conditions for semantic refinement

While semantic refinement therefore has the key property that we want from a notion of architectural refinement, establishing semantic refinement requires reasoning about the functions $\mathtt{fta}_u$, which have a complicated inductive definition that, for a given domain u, potentially ranges over the entire set of domains. It would be advantageous if refinements could be proved in a more local way. We therefore develop a number of sufficient conditions for semantic refinement that have this property.

It proves to be helpful to restrict attention to interpreted architectures satisfying a stronger version of non-conflation than that introduced above. Say that an interpreted architecture is *strongly non-conflating*, if for all $u, v, w \in D$ such that $u \xrightarrow{f} w$ with $f \neq \top$ and $v \xrightarrow{\top} w$, we have that for all sequences $\alpha, \beta \in A^*$ and actions $a, b \in A$ with $\mathtt{dom}(a) = u$ and $\mathtt{dom}(b) = v$ that $\mathtt{I}(f)(\alpha, a) \neq (\mathtt{fta}_v(\beta), b)$. Intuitively, this says that w is able to distinguish between information received from u and v, or alternately, that domain v, which is unconstrained in the information it is able to send to w, is moreover able to authenticate its messages to w. Note that the special case where $v = w$ gives that a strongly non-conflating architecture is also non-conflating.

As a first step towards a local definition of refinement, we note that it suffices to focus on the information that is "transmitted" when an action occurs, as represented by the functions $\mathtt{T}^{\mathcal{AI}}_{u,v}$. Let $\mathcal{AI}_1 = (\mathcal{A}_1, \mathcal{I}_1)$ and $\mathcal{AI}_2 = (\mathcal{A}_2, \mathcal{I}_2)$ be fully filtered interpreted extended architectures, where $\mathcal{A}_i = (D_i, \rightarrowtail_i)$, and $\mathcal{I}_i = (A, \mathtt{dom}_i, \mathtt{I}_i)$ for $i = 1, 2$. (As above, we require that the interpretations have the same set of actions A.) Formally, define a refinement mapping from $\mathcal{AI}_1$ to $\mathcal{AI}_2$ to be a function $r : D_1 \to D_2$ such that the following conditions hold.

T1 The function r is onto D_2, and $\mathtt{dom}_2 = r \circ \mathtt{dom}_1$.

T2 For all $u, v \in D_1$, $\alpha \in A^*$ and $a \in A$ with $\mathtt{dom}_1(a) = u$, if $\mathtt{T}^{\mathcal{AI}_2}_{r(u),r(v)}(\alpha, a) = \epsilon$, then $\mathtt{T}^{\mathcal{AI}_1}_{u,v}(\alpha, a) = \epsilon$.

T3 For all $u \in D_2$ and $v \in D_1$, $\alpha, \beta \in A^*$ and $a, b \in A$, if

$$\mathtt{T}^{\mathcal{AI}_2}_{\mathtt{dom}_2(a),r(v)}(\alpha, a) = \mathtt{T}^{\mathcal{AI}_2}_{\mathtt{dom}_2(b),r(v)}(\beta, b) \neq \epsilon$$

then either (i) $a = b$ and $\mathtt{dom}_2(a) \xrightarrow{\top} r(v)$ or (ii) $\mathtt{T}^{\mathcal{AI}_1}_{\mathtt{dom}_1(a),v}(\alpha, a) = \mathtt{T}^{\mathcal{AI}_1}_{\mathtt{dom}_1(b),v}(\beta, b)$.

We write $\mathcal{AI}_1 \sqsubseteq_r \mathcal{AI}_2$ when these conditions are satisfied.

Intuitively, condition T2 says that in situations where no information is permitted to flow, by the abstract architecture $\mathcal{AI}_2$, between superdomains of the concrete domains u, v, no information is permitted to flow between the domains

u, v. It can be seen to be a generalization of the refinement condition (for architectures without filter functions) considered by van der Meyden [2012]. Stated in the contrapositive, T2 says that if $\mathtt{T}^{\mathcal{AI}_1}_{u,v}(\alpha, a) \neq \epsilon$ then $\mathtt{T}^{\mathcal{AI}_2}_{r(u),r(v)}(\alpha, a) \neq \epsilon$. In particular, in the case where $u \rightarrowtail^{\top} v$, we have $\mathtt{T}^{\mathcal{AI}_1}_{u,v}(\alpha, a) = (\mathtt{fta}_u(\alpha), a) \neq \epsilon$, so $\mathtt{T}^{\mathcal{AI}_2}_{r(u),r(v)}(\alpha, a) \neq \epsilon$. This requires that there exists an edge $r(u) \rightarrowtail^{g} r(v)$.

Condition T3 is somewhat technical, but intuitively states that if it is possible for a superdomain $r(v)$ to detect the occurrence of both action a performed after α, and action b performed after β, but is not able to distinguish these actions, then it should also not be possible for the subdomain v to distinguish these actions, as stated in case (ii) of the condition. We separate out case (i) for technical reasons. It covers one situation, easily checked by inspection of the architecture, where it can be shown that v cannot distinguish the two actions (we prove this in the context of Theorem 10 below).

We now show that this more localized notion of refinement provides a sufficient condition for semantic refinement. In what follows, to lighten the notation, we drop the superscripts $\mathcal{AI}_1$ and $\mathcal{AI}_2$ from terms like $\mathtt{fta}^{\mathcal{AI}_1}_u(\alpha)$ and $\mathtt{T}^{\mathcal{AI}_2}_{\mathtt{dom}_2(a),r(u)}(\alpha, a)$, since it will always remain clear from the domains (here $u \in D_1$ or $\mathtt{dom}_2(a), r(u) \in D_2$) which architecture is intended.

Theorem 10 *Suppose that $\mathcal{AI}_1$ and $\mathcal{AI}_2$ are strongly non-conflating interpreted architectures. If $\mathcal{AI}_1 \sqsubseteq_r \mathcal{AI}_2$ then $\mathcal{AI}_1 \preceq_r \mathcal{AI}_2$.*

Proof: Suppose $\mathcal{AI}_1 \sqsubseteq_r \mathcal{AI}_2$. We need to show that for all domains $u \in D_1$ and sequences $\alpha, \beta \in A^*$, if $\mathtt{fta}_{r(u)}(\alpha) = \mathtt{fta}_{r(u)}(\beta)$ then $\mathtt{fta}_u(\alpha) = \mathtt{fta}_u(\beta)$. We show this by induction on the combined length of α and β. The case of $\alpha = \beta = \epsilon$ is trivial. Consider sequences αa and β, such that $\mathtt{fta}_{r(u)}(\alpha a) = \mathtt{fta}_{r(u)}(\beta)$, where the claim holds for shorter sequences. We consider two cases, depending on whether $\mathtt{T}_{\mathtt{dom}_2(a),r(u)}(\alpha, a) = \epsilon$.

If $\mathtt{T}_{\mathtt{dom}_2(a),r(u)}(\alpha, a) = \epsilon$, then by condition T2, we have $\mathtt{T}_{\mathtt{dom}_1(a),u}(\alpha, a) = \epsilon$. In this case, we also have, by definition, that

$$\mathtt{fta}_{r(u)}(\alpha) = \mathtt{fta}_{r(u)}(\alpha) \,\hat{}\, \mathtt{T}_{\mathtt{dom}_2(a),r(u)}(\alpha, a) = \mathtt{fta}_{r(u)}(\alpha a) = \mathtt{fta}_{r(u)}(\beta).$$

Thus, by the induction hypothesis, we have $\mathtt{fta}_u(\alpha) = \mathtt{fta}_u(\beta)$. Since $\mathtt{T}_{\mathtt{dom}_1(a),u}(\alpha, a) = \epsilon$, we obtain that

$$\mathtt{fta}_u(\alpha a) = \mathtt{fta}_u(\alpha) \,\hat{}\, \mathtt{T}_{\mathtt{dom}_1(a),u}(\alpha, a) = \mathtt{fta}_u(\alpha) = \mathtt{fta}_u(\beta),$$

as required.

Alternately, if $\mathtt{T}_{\mathtt{dom}_2(a),r(u)}(\alpha, a) \neq \epsilon$, then since

$$\mathtt{fta}_{r(u)}(\beta) = \mathtt{fta}_{r(u)}(\alpha a) = \mathtt{fta}_{r(u)}(\alpha) \,\hat{}\, \mathtt{T}_{\mathtt{dom}_2(a),r(u)}(\alpha, a),$$

it follows that β is not ϵ. We may therefore write $\beta = \beta' b$, where $b \in A$. We now have two further cases. If $\mathtt{T}_{\mathtt{dom}_2(b),r(u)}(\beta', b) = \epsilon$, then we may apply the argument above with the roles of αa and $\beta' b$ switched. Otherwise, $\mathtt{T}_{\mathtt{dom}_2(b),r(u)}(\beta', b) \neq \epsilon$. Since $\mathtt{fta}_{r(u)}(\alpha a) = \mathtt{fta}_{r(u)}(\beta' b)$ states that

$$\mathtt{fta}_{r(u)}(\alpha) \,\hat{}\, \mathtt{T}_{\mathtt{dom}_2(a),r(u)}(\alpha, a) = \mathtt{fta}_{r(u)}(\beta') \,\hat{}\, \mathtt{T}_{\mathtt{dom}_2(b),r(u)}(\beta', b)$$

and neither of the appended elements is ϵ, it follows that $\mathtt{fta}_{r(u)}(\alpha) = \mathtt{fta}_{r(u)}(\beta')$ and $\mathtt{T}_{\mathtt{dom}_2(a),r(u)}(\alpha, a) = \mathtt{T}_{\mathtt{dom}_2(b),r(u)}(\beta', b)$. By the induction hypothesis, we obtain that $\mathtt{fta}_u(\alpha) = \mathtt{fta}_u(\beta')$. Also, by condition T3 we have either (i) $a = b$ and $\mathtt{dom}_2(a) \overset{\top}{\rightarrowtail} r(u)$ or (ii) $\mathtt{T}_{\mathtt{dom}_1(a),u}(\alpha, a) = \mathtt{T}_{\mathtt{dom}_1(b),u}(\beta', b)$. We claim that (ii) also holds in case (i). It then follows, by definition, that $\mathtt{fta}_u(\alpha a) = \mathtt{fta}_u(\beta' b)$, as required.

To prove the claim, note that in case (i), we have

$$\mathtt{T}_{\mathtt{dom}_2(a),r(u)}(\alpha, a) = \mathtt{T}_{\mathtt{dom}_2(b),r(u)}(\beta', b) = \mathtt{T}_{\mathtt{dom}_2(a),r(u)}(\beta', a)$$

and since we have an edge $\mathtt{dom}_2(a) \overset{\top}{\rightarrowtail} r(u)$ this states that $(\mathtt{fta}_{\mathtt{dom}_2(a)}(\alpha), a) = (\mathtt{fta}_{\mathtt{dom}_2(a)}(\beta'), a)$. In particular, we have $\mathtt{fta}_{\mathtt{dom}_2(a)}(\alpha) = \mathtt{fta}_{\mathtt{dom}_2(a)}(\beta')$. By induction, it follows that $\mathtt{fta}_{\mathtt{dom}_1(a)}(\alpha) = \mathtt{fta}_{\mathtt{dom}_1(a)}(\beta')$. There are now several possibilities, depending on the nature of the edge, if any, from $\mathtt{dom}_1(a) = \mathtt{dom}_1(b)$ to u in $\mathcal{AI}_1$.

1. if $\mathtt{dom}_1(a) \not\rightarrowtail u$, then by definition, $\mathtt{T}_{\mathtt{dom}_1(a),u}(\alpha, a) = \epsilon = \mathtt{T}_{\mathtt{dom}_1(b),u}(\beta', b)$.

2. If $\mathtt{dom}_1(a) \overset{\top}{\rightarrowtail} u$, then $\mathtt{T}_{\mathtt{dom}_1(a),u}(\alpha, a) = (\mathtt{fta}_{\mathtt{dom}_1(a)}(\alpha), a) = (\mathtt{fta}_{\mathtt{dom}_1(b)}(\beta'), b) = \mathtt{T}_{\mathtt{dom}_1(b),u}(\beta', b)$.

3. If $\mathtt{dom}_1(a) \overset{f_1}{\rightarrowtail} u$ with $f_1 \neq \top$, then it follows from $\mathtt{fta}_{\mathtt{dom}_1(a)}(\alpha) = \mathtt{fta}_{\mathtt{dom}_1(a)}(\beta')$, using compatibility of $\mathtt{I}_1(f_1)$, that $\mathtt{T}_{\mathtt{dom}_1(a),u}(\alpha, a) = \mathtt{I}_1(f_1)(\alpha, a) = \mathtt{I}_1(f_1)(\beta', a) = \mathtt{I}_1(f_1)(\beta', b) = \mathtt{T}_{\mathtt{dom}_1(b),u}(\beta', b)$.

Thus, in all cases, we have $\mathtt{T}_{\mathtt{dom}_1(a),u}(\alpha, a) = \mathtt{T}_{\mathtt{dom}_1(b),u}(\beta', b)$. □

The refinement relation $\sqsubseteq$ is quite general and captures the essence of the relation between semantic architectural refinement and refinement of the information that flows over the edges of an architecture. However, to establish that the conditions hold may still require reasoning about the complex "global" functions $\mathtt{fta}_u$. For example, in the case where an edge $r(u) \overset{f}{\rightarrowtail} r(v)$ with $f \neq \top$ is refined by an edge $u \overset{\top}{\rightarrowtail} v$, we need to prove a relationship between the "local" function $\mathtt{I}_2(f)$, and the recursive function $\mathtt{fta}_u$. We therefore develop a more specific set of conditions that do support a more local approach to proof of a refinement, in which we need to consider only the edges not labelled by $\top$.

Define a *strict refinement mapping* from $\mathcal{AI}_1$ to $\mathcal{AI}_2$ to be a function $r : D_1 \to D_2$ such that the following conditions hold.

S1 The function r is onto D_2, and $\mathtt{dom}_2 = r \circ \mathtt{dom}_1$.

S2 For $u, v \in D_1$, if $u \overset{\top}{\rightarrowtail} v$ then $r(u) \overset{\top}{\rightarrowtail} r(v)$.

S3 For all $u, v \in D_1$, if $u \overset{f_1}{\rightarrowtail} v$ with $f_1 \neq \top$ then there exists an edge $r(u) \overset{f_2}{\rightarrowtail} r(v)$ with either $f_2 = \top$ or $f_2 \neq \top$ and for all $\alpha \in A^*$ and $a \in A$ with $\mathtt{dom}_1(a) = u$ we have $\mathtt{I}_2(f_2)(\alpha, a) = \epsilon$ implies $\mathtt{I}_1(f_1)(\alpha, a) = \epsilon$.

S4 For all $u \in D_1$ and actions $a, b \in A$, if $\mathtt{dom}_2(a) \overset{f_2}{\rightarrowtail} r(u)$ and $\mathtt{dom}_2(b) \overset{g_2}{\rightarrowtail} r(u)$ where $f_2 \neq \top$ and $g_2 \neq \top$, then either

(a) $\mathtt{dom}_1(a) \not\rightarrowtail u$ and $\mathtt{dom}_1(b) \not\rightarrowtail u$, or

(b) there exists an edge $\mathtt{dom}_1(a) \overset{f_1}{\rightarrowtail} u$ with $f_1 \neq \top$ but $\mathtt{dom}_1(b) \not\rightarrowtail u$, and for all $\alpha, \beta \in A^*$ such that $\mathtt{I}_2(f_2)(\alpha, a) = \mathtt{I}_2(g_2)(\beta, b) \neq \epsilon$ we have $\mathtt{I}_1(f_1)(\alpha, a) = \epsilon$, or

(c) $\mathtt{dom}_1(a) \not\rightarrowtail u$ but there exists an edge $\mathtt{dom}_1(b) \overset{g_1}{\rightarrowtail} u$ with $g_1 \neq \top$, and for all $\alpha, \beta \in A^*$ such that $\mathtt{I}_2(f_2)(\alpha, a) = \mathtt{I}_2(g_2)(\beta, b) \neq \epsilon$ we have $\mathtt{I}_1(g_1)(\beta, b) = \epsilon$, or

(d) there exist edges $\mathtt{dom}_1(a) \overset{f_1}{\rightarrowtail} v$ and $\mathtt{dom}_2(b) \overset{g_1}{\rightarrowtail} v$ with $f_1 \neq \top$ and $g_1 \neq \top$, and for all $\alpha, \beta \in A^*$, such that $\mathtt{I}_2(f_2)(\alpha, a) = \mathtt{I}_2(g_2)(\beta, b) \neq \epsilon$ we have $\mathtt{I}_1(f_1)(\alpha, a) = \mathtt{I}_1(g_1)(\beta, b)$.

Note that in the case where all edges in the two architectures are labelled by $\top$, conditions S1-2 amount to the notion of architectural refinement of van der Meyden [2012]. For edges labelled with non-$\top$ labels, Condition S3 states that whenever information is not permitted to flow between two superdomains, it is not permitted to flow between their subdomains.

Condition S4, like condition T3, is intended to capture that if a superdomain cannot distinguish two actions a and b, then neither can its subdomains. However, S4 restricts the statement of this property to filtered edges, i.e., edges not labelled $\top$. The essence of the property is most easily visible in condition S4(d), which corresponds to a situation where both $\mathtt{dom}_1(a)$ and $\mathtt{dom}_1(b)$ have filtered edges to a subdomain. The remaining cases deal with all the other possible configurations that are consistent with S2.

Note that cases (b) and (c) of S4 say that where a superdomain $r(u)$ is not able to distinguish the actions a and b, it is not permitted for the subdomain u to receive information flow as a result of one action but not the other. Intuitively, this would imply that the subdomain can deduce which action occurred, giving it more information than its superdomain, whereas our intuition for refinement is that it should not increase the amount of information flow.

We remark that S4 does not need to consider the situation where $\mathtt{dom}_2(a) \overset{f_2}{\rightarrowtail} r(u)$ with $f_2 \neq \top$ and $\mathtt{dom}_2(b) \overset{\top}{\rightarrowtail} r(u)$ because in strongly non-conflating architectures, it is impossible that $\mathtt{T}_{\mathtt{dom}_2(a), r(u)}(\alpha, a) = \mathtt{T}_{\mathtt{dom}_2(b), r(u)}(\beta, b)$.

Theorem 11 *If $\mathcal{AI}_2$ is strongly non-conflating and $\mathcal{AI}_1$ strictly refines $\mathcal{AI}_2$ by function r, then $\mathcal{AI}_1 \sqsubseteq_r \mathcal{AI}_2$.*

Proof: Assume that $\mathcal{AI}_1$ and $\mathcal{AI}_2$ are strongly non-conflating and that conditions S1–4 hold. We prove conditions T1-3. T1 is trivially identical to S1.

For condition T2, we consider that contrapositive. Suppose that $u, v \in D_1$ and $\alpha \in A^*$ and $a \in A$ with $\mathtt{dom}_1(a) = u$ and $\mathtt{T}_{u,v}(\alpha, a) \neq \epsilon$. We need to prove that $\mathtt{T}_{r(u), r(v)}(\alpha, a) \neq \epsilon$. From $\mathtt{T}_{u,v}(\alpha, a) \neq \epsilon$ we obtain that there exists an edge $u \overset{f_1}{\rightarrowtail} v$. If $f_1 = \top$, then by S2, we have that $r(u) \overset{\top}{\rightarrowtail} r(v)$, and since $\mathtt{dom}_2(a) =$

$r(\mathtt{dom}_1(a)) = r(u)$ by S1, this implies that $\mathtt{T}_{r(u),r(v)}(\alpha, a) = (\mathtt{fta}_{r(u)}(\alpha), a) \neq \epsilon$. Alternately, if $f_1 \neq \top$, then by S3, there exists an edge $r(u) \overset{f_2}{\rightarrowtail} r(v)$, with either $f_2 = \top$ or $f_2 \neq \top$ and $\mathtt{I}_2(f_2)(\alpha, a) = \epsilon$ implies $\mathtt{I}_1(f_1)(\alpha, a) = \epsilon$. In the case $f_2 = \top$ we argue exactly as above. In case $f_2 \neq \top$, since we have $\mathtt{I}_1(f_1)(\alpha, a) = \mathtt{T}_{u,v}(\alpha a) \neq \epsilon$, we obtain $\epsilon \neq \mathtt{I}_2(f_2)(\alpha, a) = \mathtt{T}_{r(u),r(v)}(\alpha, a)$, again as required.

For condition T3, suppose $\mathtt{T}_{\mathtt{dom}_2(a),r(v)}(\alpha, a) = \mathtt{T}_{\mathtt{dom}_2(b),r(v)}(\beta, b) \neq \epsilon$. We have to show that either (i) $a = b$ and $\mathtt{dom}_2(a) \overset{\top}{\rightarrowtail} r(v)$, or (ii) $\mathtt{T}_{\mathtt{dom}_1(a),v}(\alpha, a) = \mathtt{T}_{\mathtt{dom}_1(b),v}(\beta, b)$. From $\mathtt{T}_{\mathtt{dom}_2(a),r(v)}(\alpha, a) = \mathtt{T}_{\mathtt{dom}_2(b),r(v)}(\beta, b) \neq \epsilon$ it follows that there exist edges $\mathtt{dom}_2(a) \overset{f_2}{\rightarrowtail} r(v)$ and $\mathtt{dom}_2(a) \overset{g_2}{\rightarrowtail} r(v)$. We consider several cases, depending on whether these edges are labelled $\top$ or not.

If $f_2 = g_2 = \top$, then $\mathtt{T}_{\mathtt{dom}_2(a),r(v)}(\alpha, a) = (\mathtt{fta}_{\mathtt{dom}_2(a)}(\alpha), a)$ and $\mathtt{T}_{\mathtt{dom}_2(b),r(v)}(\beta, b) = (\mathtt{fta}_{\mathtt{dom}_2(b)}(\beta), b)$, and it follows that $a = b$ and we have (i).

The case that $f_2 = \top$ and $g_2 \neq \top$ is not possible, by the assumption that the architecture $\mathcal{AI}_2$ is strongly non-conflating.

If $f_2 \neq \top$ and $g_2 \neq \top$, then by S4, we have one of four possibilities.

(a) $\mathtt{dom}_1(a) \not\rightarrowtail v$ and $\mathtt{dom}_1(b) \not\rightarrowtail v$. In this case, $\mathtt{T}_{\mathtt{dom}_1(a),v}(\alpha a) = \epsilon = \mathtt{T}_{\mathtt{dom}_1(b),v}(\beta b)$.

(b) There exists an edge $\mathtt{dom}_1(a) \overset{f_1}{\rightarrowtail} v$ with $f_1 \neq \top$ but $\mathtt{dom}_1(b) \not\rightarrowtail v$, and $\mathtt{I}_2(f_2)(\alpha, a) = \mathtt{I}_2(g_2)(\beta, b) \neq \epsilon$ implies $\mathtt{I}_1(f_1)(\alpha, a) = \epsilon$. In this case, since

$$\mathtt{I}_2(f_2)(\alpha, a) = \mathtt{T}_{\mathtt{dom}_2(a),r(v)}(\alpha, a) = \mathtt{T}_{\mathtt{dom}_2(b),r(v)}(\beta, b) = \mathtt{I}_2(g_2)(\beta, b) \neq \epsilon,$$

we obtain that $\mathtt{I}_1(f_1)(\alpha, a) = \epsilon$. Thus, $\mathtt{T}_{\mathtt{dom}_1(a),v}(\alpha, a) = \mathtt{I}_1(\alpha, a) = \epsilon = \mathtt{T}_{\mathtt{dom}_1(b),v}(\beta, b)$.

(c) This case is identical to case (b) with the roles of αa and βb reversed.

(d) There exists edges $\mathtt{dom}_1(a) \overset{f_1}{\rightarrowtail} v$ and $\mathtt{dom}_1(b) \overset{g_1}{\rightarrowtail} v$ with $f_1 \neq \top$ and $g_1 \neq \top$ and $\mathtt{I}_2(f_2)(\alpha, a) = \mathtt{I}_2(g_2)(\beta, b) \neq \epsilon$ implies $\mathtt{I}_1(f_1)(\alpha, a) = \mathtt{I}_1(g_1)(\beta, b)$. Since $\mathtt{I}_2(f_2)(\alpha, a) = \mathtt{T}_{\mathtt{dom}_2(a),r(v)}(\alpha, a) = \mathtt{T}_{\mathtt{dom}_2(b),r(v)}(\beta, b) = \mathtt{I}_2(g_2)(\beta, b) \neq \epsilon$ we obtain $\mathtt{I}_1(f_1)(\alpha, a) = \mathtt{I}_1(g_1)(\beta, b)$, which is identical to $\mathtt{T}_{\mathtt{dom}_1(a),v}(\alpha, a) = \mathtt{T}_{\mathtt{dom}_1(b),v}(\beta, b)$.

Thus, in each case, we have $\mathtt{T}_{\mathtt{dom}_1(a),v}(\alpha, a) = \mathtt{T}_{\mathtt{dom}_1(b),v}(\beta, b)$, as required. □

We remark that although condition S4 is somewhat complex, under a reasonable assumption it can be replaced by the following much simpler condition S5, which states more transparently that at least as much information is permitted to flow along abstract edges as is permitted to flow along any corresponding concrete edges.

S5 For $\alpha, \beta \in A^*$ and $a, b \in A$ and $v \in D_1$ such that $\mathtt{dom}_1(a) = \mathtt{dom}_1(b)$, and $\mathtt{dom}_1(a) \overset{f_1}{\rightarrowtail}_1 v$ and $\mathtt{dom}_2(a) \overset{f_2}{\rightarrowtail}_2 r(v)$, where $f_1 \neq \top$ and $f_2 \neq \top$, if $\mathtt{I}_2(f_2)(\alpha, a) = \mathtt{I}_2(f_2)(\beta, b) \neq \epsilon$, then $\mathtt{I}_1(f_1)(\alpha, a) = \mathtt{I}_1(f_1)(\beta, b)$.

Say that *messages are source-identifying with respect to r*, if there exists a function $\mathcal{S}$ mapping the union of the ranges of the filter functions $\mathtt{I}_2(f)$ of $\mathcal{AI}_2$ to the set of domains D_1 of $\mathcal{AI}_1$, such that for all $\alpha \in A^*$ and $a \in A$, if $r(\mathtt{dom}_1(a)) \overset{f_2}{\rightarrowtail}_2 w$ and $\mathtt{I}_2(f_2)(\alpha, a) \neq \epsilon$, then $\mathcal{S}(\mathtt{I}_2(f_2)(\alpha, a)) = \mathtt{dom}_1(a)$. One example of when this condition can be met is when $\mathtt{I}_2(f_2)(\alpha, a) = (g(\alpha), a)$ for some function g, since then the function $\mathcal{S}((x, a)) = \mathtt{dom}_1(a)$ obviously has the required property.

Lemma 6 *If messages are source-identifying with respect to r and r satisfies conditions S1, S2 and S5 then r satisfies condition S4.*

Proof: Suppose $\mathtt{dom}_2(a) \overset{f_2}{\rightarrowtail}_2 r(v)$, $\mathtt{dom}_2(b) \overset{g_2}{\rightarrowtail}_2 r(v)$ with $f_2 \neq \top$ and $g_2 \neq \top$. Let $u \in D_1$. By condition S2, we cannot have $\mathtt{dom}_1(a) \overset{\top}{\rightarrowtail} u$ or $\mathtt{dom}_1(b) \overset{\top}{\rightarrowtail} u$. This leaves four possibilities, depending on whether there is an non-$\top$-labelled edge between $\mathtt{dom}_1(a)$ or $\mathtt{dom}_1(b)$ and u or not. In case there are no such edges, we have condition S4(a). We consider the three other possibilities. For each case, observe that if there exists $\alpha, \beta \in A^*$ with $\mathtt{I}_2(f_2)(\alpha, a) = \mathtt{I}_2(g_2)(\beta, b) \neq \epsilon$, then $\mathtt{dom}_1(a) = \mathcal{S}(\mathtt{I}(f_2)(\alpha, a)) = \mathcal{S}(\mathtt{I}(g_2)(\beta, b)) = \mathtt{dom}_1(b)$, by the assumption that messages are source-identifying with respect to r. Conversely, if $\mathtt{dom}_1(a) \neq \mathtt{dom}_1(b)$, then $\mathtt{I}_2(f_2)(\alpha, a) = \mathtt{I}_2(g_2)(\beta, b) \neq \epsilon$ cannot be satisfied.

- Suppose there exists an edge $\mathtt{dom}_1(a) \overset{f_1}{\rightarrowtail} u$ with $f_1 \neq \top$ and $\mathtt{dom}_1(b) \not\rightarrowtail u$. In this case, we must have $\mathtt{dom}_1(a) \neq \mathtt{dom}_1(b)$, so by the observation above, $\mathtt{I}_2(f_2)(\alpha, a) = \mathtt{I}_2(g_2)(\beta, b) \neq \epsilon$ cannot be satisfied. Thus, condition S4(b) is vacuously satisfied.

- Suppose $\mathtt{dom}_1(b) \not\rightarrowtail u$ but there exists an edge $\mathtt{dom}_1(a) \overset{g_1}{\rightarrowtail} u$ with $g_1 \neq \top$. This case is similar to the previous one, yielding S4(c).

- Suppose there exist edges $\mathtt{dom}_1(a) \overset{f_1}{\rightarrowtail} u$ and $\mathtt{dom}_1(a) \overset{g_1}{\rightarrowtail} u$ with $f_1, g_1 \neq \top$. If $\mathtt{I}_2(f_2)(\alpha, a) = \mathtt{I}_2(g_2)(\beta, b) \neq \epsilon$, then as observed above, we have $\mathtt{dom}_1(a) = \mathtt{dom}_1(b)$, and hence also $f_1 = g_1$ by the fact that there is at most one edge between any two domains. Thus, by S5, we have $\mathtt{I}(f_1)(\alpha, a) = \mathtt{I}(f_1)(\beta, b) = \mathtt{I}(g_1)(\beta, b)$. This establishes condition S4(d).

Thus, in all cases, we have proved condition S4. □

5.1.2 Properties of architectural refinement

Architectural refinement provides a design methodology in which we may prove certain security properties at a high level of abstraction, and preserve the validity of those properties as details of the architectural design are specified. Since an interpreted architecture $\mathcal{AI}$ specifies a set of actions A, any interpretation function π mapping atomic propositions to subsets of A^* can be treated as an interpretation for any system M that is FTA-compliant with $\mathcal{AI}$. For a formula ϕ, we write $\mathcal{AI}, \pi \models \phi$ if $M, \pi \models \phi$ for all systems M that are FTA-compliant with $\mathcal{AI}$.

Recall that $\mathcal{AI}_1 \preceq_r \mathcal{AI}_2$ entails that $\mathcal{AI}_1$ and $\mathcal{AI}_2$ have the same set of actions A. Since interpretations π map each propositional constant to a subset of A^*, the two interpreted architectures also have the same interpretations. The following result shows that a property that has been shown to follow from compliance with $\mathcal{AI}_2$ can be translated to a property that follows from compliance with $\mathcal{AI}_1$.

Theorem 12 *Suppose $\mathcal{AI}_1$ and $\mathcal{AI}_2$ are interpreted extended architectures such that $\mathcal{AI}_1 \preceq_r \mathcal{AI}_2$. Let π be an interpretation for these architectures. Suppose ϕ is a formula for the domains of $\mathcal{AI}_2$, and $\mathcal{AI}_2, \pi \vDash \phi$. Then $\mathcal{AI}_1, \pi \vDash r^{-1}(\phi)$.*

Proof: Suppose that system M is FTA-compliant with $\mathcal{AI}_1$. Since $\mathcal{AI}_1 \preceq_r \mathcal{AI}_2$, it follows by Theorem 9 that $r(M)$ is FTA-compliant with $\mathcal{AI}_2$. Hence $r(M), \pi \vDash \phi$. By Theorem 1 we obtain that $M, \pi \models r^{-1}(\phi)$. □

The level of abstraction used to model a system may affect the success or efficiency of a proof of a given global security property. Theorem 12 can facilitate proofs of security properties that are preserved under refinement: simpler, more abstract, architectures can be used to reason about the security property, and preservation under refinement ensures that the security property will hold of systems satisfying a more refined architecture. We demonstrate this in Sections 5.3 and 5.4 below, using more abstract architectures to prove security results about the Hinke-Schaefer architecture and a refinement of the downgrader architecture.

5.2 Specification Refinement

We can also define refinement at the level of architectural specifications. Suppose that $\mathcal{A}_1 = (D_1, \rightarrowtail_1)$ and $\mathcal{A}_2 = (D_2, \rightarrowtail_2)$ are architectures, and r is a function mapping D_1 onto D_2. If $\mathcal{C}_1$ and $\mathcal{C}_2$ are architectural specifications for $\mathcal{A}_1$ and $\mathcal{A}_2$, respectively, then we write $(\mathcal{A}_1, \mathcal{C}_1) \preceq_r (\mathcal{A}_2, \mathcal{C}_2)$ when for all interpretations $\mathcal{I}_1 \in \mathcal{C}_1$ there exists an interpretation $\mathcal{I}_2 \in \mathcal{C}_2$ such that $(\mathcal{A}_1, \mathcal{I}_1) \preceq_r (\mathcal{A}_2, \mathcal{I}_2)$.

Architectural specifications may allow multiple architectural interpretations, and these may allow the set of actions A in the systems being specified to vary. In order to interpret the logic of knowledge, we need an interpretation π that maps each propositional constant to a subset of A^*. When A varies, we need π to vary correspondingly. To ensure proper coordination between interpretations π and sets of actions A, we work with sets $\mathcal{M}$ of *interpreted systems*, i.e., pairs (M, π) where M is a system and π is an interpretation. Such a set $\mathcal{M}$ may express weak conditions such as "$\pi(p)$ is a G-dependent proposition". We write $\mathcal{M}, (\mathcal{A}, \mathcal{C}) \models \phi$ when $M, \pi \models \phi$ for all $(M, \pi) \in \mathcal{M}$ such that M FTA-complies with $(\mathcal{A}, \mathcal{C})$. We also write $r(\mathcal{M})$ for $\{(r(M), \pi) \mid (M, \pi) \in \mathcal{M}\}$.

Theorem 13 *Let $\mathcal{M}_1, \mathcal{M}_2$ be sets of interpreted systems for architectures $\mathcal{A}_1$, $\mathcal{A}_2$, respectively, and let ϕ be a formula for the domains of $\mathcal{A}_2$. Suppose that $(\mathcal{A}_1, \mathcal{C}_1) \preceq_r (\mathcal{A}_2, \mathcal{C}_2)$ and $r(\mathcal{M}_1) \subseteq \mathcal{M}_2$ and $\mathcal{M}_2, (\mathcal{A}_2, \mathcal{C}_2) \models \phi$. Then we have $\mathcal{M}_1, (\mathcal{A}_1, \mathcal{C}_1) \models r^{-1}(\phi)$.*

Proof: Let $(M, \pi) \in \mathcal{M}_1$ and suppose M is FTA-compliant with $(\mathcal{A}_1, \mathcal{C}_1)$. We need to show that $M, \pi \models r^{-1}(\phi)$. By definition, there exists an interpretation $\mathcal{I}_1$ such that M is FTA-compliant with $(\mathcal{A}_1, \mathcal{I}_1)$. Since $(\mathcal{A}_1, \mathcal{C}_1) \preceq_r (\mathcal{A}_2, \mathcal{C}_2)$, it follows that there exists an interpretation $\mathcal{I}_2$ such that $(\mathcal{A}_1, \mathcal{I}_1) \preceq_r (\mathcal{A}_2, \mathcal{I}_2)$. By Theorem 9, it follows that $r(M)$ is FTA-compliant with $(\mathcal{A}_2, \mathcal{I}_2)$, and hence also with $(\mathcal{A}_2, \mathcal{C}_2)$. Moreover, since $r(\mathcal{M}_1) \subseteq \mathcal{M}_2$, we have $(r(M), \pi) \in \mathcal{M}_2$. Thus $r(M), \pi \models \phi$. By Theorem 1 we obtain that $M, \pi \models r^{-1}(\phi)$, as required. □

We may similarly define $(\mathcal{A}_1, \mathcal{C}_1) \leq_r (\mathcal{A}_2, \mathcal{C}_2)$ to hold when for all interpretations $\mathcal{I}_1 \in \mathcal{C}_1$ there exists an interpretation $\mathcal{I}_2 \in \mathcal{C}_2$ such that $(\mathcal{A}_1, \mathcal{I}_1) \leq_r (\mathcal{A}_2, \mathcal{I}_2)$. Additionally, we may define $(\mathcal{A}_1, \mathcal{C}_1)$ to be a strict refinement of $(\mathcal{A}_2, \mathcal{C}_2)$ by function r when for all interpretations $\mathcal{I}_1 \in \mathcal{C}_1$ there exists an interpretation $\mathcal{I}_2 \in \mathcal{C}_2$ such that $(\mathcal{A}_1, \mathcal{I}_1)$ is a strict refinement of $(\mathcal{A}_2, \mathcal{I}_2)$.

Theorem 14 *If $(\mathcal{A}_1, \mathcal{C}_1)$ is a strict refinement of $(\mathcal{A}_2, \mathcal{C}_2)$ by function r then $(\mathcal{A}_1, \mathcal{C}_1) \sqsubseteq_r (\mathcal{A}_2, \mathcal{C}_2)$, and if $(\mathcal{A}_1, \mathcal{C}_1) \sqsubseteq_r (\mathcal{A}_2, \mathcal{C}_2)$ then $(\mathcal{A}_1, \mathcal{C}_1) \preceq_r (\mathcal{A}_2, \mathcal{C}_2)$.*

Proof: Straightforward using Theorem 10 and Theorem 11. □

5.3 Example: Hinke-Schaefer

Since the Hinke-Schaefer architecture $\mathcal{HS}$ refines architecture $\mathcal{HL}$, we can apply the information security result for $\mathcal{HL}$, Theorem 2, to $\mathcal{HS}$: since domain L never knows any H-dependent proposition, the domains L_{user}, L_{DBMS}, and L_F never know any u-dependent proposition, for $u \in \{H_{user}, H_{DBMS}, H_F\}$. This information security property was stated as Theorem 3, in Section 2.2. We give a simple proof for it here.

Proof of Theorem 3: Follows easily from Theorem 2, Lemma 3, and Theorem 12, since $\mathcal{HS} \leq_r \mathcal{HL}$ for refinement function r such that $r(L_{user}) = r(L_{DBMS}) = r(L_F) = L$ and $r(H_{user}) = r(H_{DBMS}) = r(H_F) = H$. □

Thus, we were able to prove an information security property about $\mathcal{HS}$ by proving an appropriate policy in the much simpler architecture $\mathcal{HL}$.

5.4 Example: Downgrader

In the architecture $\mathcal{DG}$ of Section 4.3, there is a filter function on the edge from downgrader D to Low security domain L, specifying that D should not release to L any information about C. This makes D a trusted component in the system: in any implementation, we would need to verify that D correctly enforces this information flow constraint. However, information about C may become comingled with information about P in the data store H, so it is not immediately clear how D could, on its own, guarantee enforcement of the constraint. Thus, it seems that the architecture implies constraints on other components. One approach that could be pursued to implement this architecture is to ensure that the data store H maintains secure provenance information, which D can use to check that information being released is not tainted with information from

C. In this section, we pursue another approach, which is to move the trust boundary in such a way that D is prevented from obtaining information about C. We develop an architecture $\mathcal{DGR}$ that has this property, and show it to be a refinement of $\mathcal{DG}$. It follows that any system compliant with $\mathcal{DGR}$ is also complaint with $\mathcal{DG}$. Since it is much clearer how architecture $\mathcal{DGR}$ could be implemented with just local verification of its trusted components, this moves us closer to a practical implementation of $\mathcal{DG}$.

The following diagram shows architecture $\mathcal{DGR}$, with grouping of its domains (indicated by dashed rectangles) indicating a refinement mapping r to the architecture $\mathcal{DG}$. In $\mathcal{DGR}$, domain H of $\mathcal{DG}$ is decomposed into two domains H_C and H_P corresponding respectively to data stores of information about C and P. Domain P is decomposed into two domains T and U, corresponding to trusted and untrusted users within this domain. Thus, the mapping r from the domains of $\mathcal{DGR}$ to the domains of $\mathcal{DG}$ is given by $r(H_P) = r(H_C) = H$, $f(T) = r(U) = P$ and $r(u) = u$ for $u \in \{C, D, L\}$.

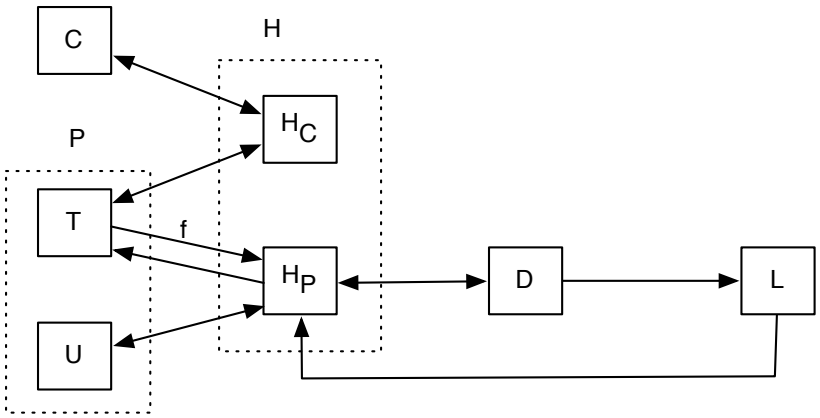

In $\mathcal{DGR}$, T is a trusted component, since this domain is required to enforce an information flow constraint, represented by the edge $T \overset{f}{\rightarrowtail} H_P$. On the other hand, note that whereas in $\mathcal{DG}$, domain D is trusted, because of the edge $D \overset{rel}{\rightarrowtail} L$, in $\mathcal{DGR}$ the edge from D to L is labelled $\top$, so D is no longer a trusted component in $\mathcal{DGR}$.

We define architectural specification $\mathcal{C}_{\mathcal{DGR}}$ so that $(A, \mathtt{dom}, \mathtt{I}) \in \mathcal{C}_{\mathcal{DGR}}$ if and only if $(A, \mathtt{dom}, \mathtt{I})$ is an interpretation of $\mathcal{DGR}$ such that $\mathtt{I}(f)$ is defined by $\mathtt{I}(f)(\alpha, a) = a$ for all $\alpha \in A^*$ and $a \in A$ with $\mathtt{dom}(a) = T$, and $\mathtt{I}(f)(\alpha, a) = \epsilon$ otherwise. Intuitively, this states that f permits information about T actions to flow from T to H_P, but nothing more. It is trivial to check that this interpretation is $\mathtt{fta}_T$-compatible.

Theorem 15 $(\mathcal{DGR}, \mathcal{C}_{\mathcal{DGR}}) \preceq_r (\mathcal{DG}, \mathcal{C}_{\mathcal{DG}})$

Proof: We need to work with three distinct architectures in this proof. To distinguish them, we use $\rightarrowtail_1$ to refer to the information flow policy of $\mathcal{DGR}$,

$\rightarrowtail_2$ to refer to the information flow policy of $\mathcal{DG}$, and $\rightarrowtail'$ to refer to the policy used in the definition of $\mathtt{I}(rel)$ (see Section 4.3). Let $\mathcal{AI}_1 = (\mathcal{DGR}, \mathcal{I}_1)$ where $\mathcal{I}_1 = (A, \mathtt{dom}_1, \mathtt{I}_1) \in \mathcal{C}_{\mathcal{DGR}}$. We need to show that there exists $\mathcal{I}_2 \in \mathcal{C}_{\mathcal{DGR}}$ with $\mathcal{AI}_1 \preceq \mathcal{AI}_2$. In fact, there exists a unique architectural interpretation $\mathcal{I}_2 = (A, \mathtt{dom}_2, \mathtt{I}_2) \in \mathcal{C}_{\mathcal{DGR}}$ with the same set A of actions and $\mathtt{dom}_2 = r \circ \mathtt{dom}_1$, so let $\mathcal{I}_2$ be this interpretation. Recall from Section 4.3 that $\mathtt{I}_2$ is defined by reference to the information flow policy $\rightarrowtail'$ once A and $\mathtt{dom}_2$ are fixed, by $\mathtt{I}_2(rel)(\alpha, a) = (\mathtt{fta}_D \rightarrowtail'(\alpha), a)$ when $\alpha \in A^*$ and $\mathtt{dom}(a) = D$.

We show that refinement mapping r satisfies conditions T1-T3. (Note that we cannot apply strict refinement, because $r(D) = D$, $r(L) = L$ and the edges $D \rightarrowtail^{\top}_1 L$ and $D \rightarrowtail^{rel}_2 L$ violate S2.) T1 is immediate from the definitions above. For T2, note the only ways that we could have $\mathtt{T}^{\mathcal{AI}_2}_{\mathtt{dom}_2(a), r(v)}(\alpha, a) = \epsilon$ for $\alpha \in A^*$ and $a \in A$ is when $\mathtt{dom}_2(a) \not\rightarrowtail_2 r(v)$ or $\mathtt{dom}_2(a) = D$ and $r(v) = L$ and $\mathtt{I}(rel)(\alpha, a) = \epsilon$. In case $\mathtt{dom}_2(a) \not\rightarrowtail_2 r(v)$, we also have $\mathtt{dom}_1(a) \not\rightarrowtail_1 v$, so also $\mathtt{T}^{\mathcal{AI}_1}_{\mathtt{dom}_1(a), v}(\alpha, a) = \epsilon$. On the other hand, the situation $\mathtt{dom}_2(a) = D$ and $r(v) = L$ and $\mathtt{I}(rel)(\alpha, a) = \epsilon$ is impossible, since $\mathtt{I}_2(rel)(\alpha, a) = (\mathtt{fta}_D(\alpha), a) \neq \epsilon$. Thus T2 holds.

For T3, we need to show that for all domains u of $\mathcal{DGR}$ and v of $\mathcal{DG}$, $\alpha \in A^*$ and $a, b \in A$, if $\mathtt{T}^{\mathcal{AI}_2}_{\mathtt{dom}_2(a), r(v)}(\alpha, a) = \mathtt{T}^{\mathcal{AI}_2}_{\mathtt{dom}_2(b), r(v)}(\beta, b) \neq \epsilon$ then either (i) $a = b$ and $\mathtt{dom}_2(a) \rightarrowtail^{\top}_2 r(v)$ or (ii) $\mathtt{T}^{\mathcal{AI}_1}_{\mathtt{dom}_1(a), v}(\alpha, a) = \mathtt{T}^{\mathcal{AI}_1}_{\mathtt{dom}_1(b), v}(\beta, b)$.

Suppose first that $r(v) \in \{C, P, H, D\}$. Note that the only incoming edges for these domains in $\mathcal{DG}$ are labelled $\top$. Hence, if $\mathtt{T}^{\mathcal{AI}_2}_{\mathtt{dom}_2(a), r(v)}(\alpha, a) = \mathtt{T}^{\mathcal{AI}_2}_{\mathtt{dom}_2(b), r(v)}(\beta, b) \neq \epsilon$ it follows that $\mathtt{dom}_2(a) \rightarrowtail^{\top}_2 r(v)$ and $\mathtt{dom}_2(b) \rightarrowtail^{\top}_2 r(v)$ and $(\mathtt{fta}_{\mathtt{dom}_2(a)}(\alpha), a) = (\mathtt{fta}_{\mathtt{dom}_2(b)}(\beta), b)$. Thus $a = b$, and we have (i).

Alternately, suppose that $r(v) = L$. If $\mathtt{dom}_2(a) = \mathtt{dom}_2(b) = L$ then the previous argument also applies. If $\mathtt{dom}_2(a) = L$ and $\mathtt{dom}_2(b) = D$ (or *vice versa*), then $\mathtt{T}^{\mathcal{AI}_2}_{\mathtt{dom}_2(a), r(v)}(\alpha, a) = \mathtt{T}^{\mathcal{AI}_2}_{\mathtt{dom}_2(b), r(v)}(\beta, b)$ states $(\mathtt{fta}^{\rightarrowtail_2}_{\mathtt{dom}_2(a)}(\alpha), a) = (\mathtt{fta}^{\rightarrowtail'}_{\mathtt{dom}_2(b)}(\beta), b)$, which is impossible since we cannot have $a = b$ when these actions are in different domains.

Thus, the only case remaining to be considered is when $r(v) = L$ and $\mathtt{dom}_2(a) = \mathtt{dom}_2(b) = D$. Here we have $\mathtt{T}^{\mathcal{AI}_2}_{\mathtt{dom}_2(a), r(v)}(\alpha, a) = (\mathtt{fta}^{\rightarrowtail'}_D(\alpha), a)$ and $\mathtt{T}^{\mathcal{AI}_2}_{\mathtt{dom}_2(b), r(v)}(\beta, b) = (\mathtt{fta}^{\rightarrowtail'}_D(\beta), b)$, so equality of these terms implies that $a = b$ and $\mathtt{fta}^{\rightarrowtail'}_D(\alpha) = \mathtt{fta}^{\rightarrowtail'}_D(\beta)$. It is not the case that $D \rightarrowtail^{\top}_2 L$, so we need to establish that $\mathtt{T}^{\mathcal{AI}_2}_{\mathtt{dom}_2(a), r(v)}(\alpha, a) = \mathtt{T}^{\mathcal{AI}_2}_{\mathtt{dom}_2(b), r(v)}(\beta, b)$, which, in light of the edge $D \rightarrowtail^{\top}_1 L$, amounts to $(\mathtt{fta}^{\rightarrowtail^1}_D(\alpha), a) = (\mathtt{fta}^{\rightarrowtail^1}_D(\beta), b)$. As we already have $a = b$, it suffices to show $\mathtt{fta}^{\rightarrowtail^1}_D(\alpha) = \mathtt{fta}^{\rightarrowtail^1}_D(\beta)$.

To obtain the required result, we claim that for all $\alpha, \beta \in A^*$, we have that $\mathtt{fta}^{\rightarrowtail'}_D(\alpha) = \mathtt{fta}^{\rightarrowtail'}_D(\beta)$ implies $\mathtt{fta}^{\rightarrowtail^1}_D(\alpha) = \mathtt{fta}^{\rightarrowtail^1}_D(\beta)$. For this, we prove that there exist functions F_u for $u \in \{T, U, H_U, D, L\}$ such that $F_u(\mathtt{fta}^{\rightarrowtail'}_{r(u)}(\alpha)) = \mathtt{fta}^{\rightarrowtail^1}_u(\alpha)$. The claim is then immediate from the case $u = D$. We define the

F_u inductively by $F_u(\epsilon) = \epsilon$ and

$$F_u(\sigma \hat{} (\delta, a)) = \begin{cases} F_u(\sigma) \hat{} (F_{\mathtt{dom}_1(a)}(\delta), a) & \text{if } \mathtt{dom}_1(a) \overset{\top}{\rightarrowtail}_1 u \\ F_u(\sigma) \hat{} a & \text{if } \mathtt{dom}_1(a) = T \text{ and } u = H_U \\ F_u(\sigma) & \text{if } \mathtt{dom}_1(a) \not\rightarrowtail_1 u \ . \end{cases}$$

Note that in case $\mathtt{dom}_1(a) \overset{\top}{\rightarrowtail}_1 u$ and $u \in \{U, H_U, D, L\}$, we must have $\mathtt{dom}_1(a) \in \{U, H_U, D, L\}$, so the recursion in the first case is well defined.

We prove by induction on $\alpha \in A^*$ that $F_u(\mathtt{fta}^{\rightarrowtail'}_{r(u)}(\alpha)) = \mathtt{fta}^{\rightarrowtail_1}_u(\alpha)$ for $u \in \{U, H_U, D, L\}$. The base case of $\alpha = \epsilon$ is trivial. Consider αa, where the statement holds for α. There are several possibilities:

Case 1: $\mathtt{dom}_1(a) \overset{\top}{\rightarrowtail}_1 u$. In this case, $\mathtt{dom}_2(a) \overset{\top}{\rightarrowtail}' r(u)$. Thus,

$$\begin{aligned} F_u(\mathtt{fta}^{\rightarrowtail'}_{r(u)}(\alpha a)) &= F_u(\mathtt{fta}^{\rightarrowtail'}_{r(u)}(\alpha) \hat{} (\mathtt{fta}^{\rightarrowtail'}_{\mathtt{dom}_2(a)}(\alpha), a)) \\ &= F_u(\mathtt{fta}^{\rightarrowtail'}_{r(u)}(\alpha)) \hat{} (F_u(\mathtt{fta}^{\rightarrowtail'}_{\mathtt{dom}_2(a)}(\alpha)), a) \\ &= \mathtt{fta}^{\rightarrowtail_1}_u(\alpha)) \hat{} (\mathtt{fta}^{\rightarrowtail_1}_{\mathtt{dom}_1(a)}(\alpha)), a) && \text{by induction} \\ &= \mathtt{fta}^{\rightarrowtail_1}_u(\alpha a). \end{aligned}$$

Case 2: $\mathtt{dom}_1(a) = T$ and $u = H_U$. In this case, $\mathtt{dom}_1(a) \overset{f}{\rightarrowtail}_1 r(u)$ and $\mathtt{dom}_2(a) = P$, $r(u) = H$, and $P \overset{\top}{\rightarrowtail}' H$. Thus,

$$\begin{aligned} F_u(\mathtt{fta}^{\rightarrowtail'}_{r(u)}(\alpha a)) &= F_u(\mathtt{fta}^{\rightarrowtail'}_{r(u)}(\alpha) \hat{} (\mathtt{fta}^{\rightarrowtail'}_{\mathtt{dom}_2(a)}(\alpha), a)) \\ &= F_u(\mathtt{fta}^{\rightarrowtail'}_{r(u)}(\alpha)) \hat{} a \\ &= \mathtt{fta}^{\rightarrowtail_1}_u(\alpha)) \hat{} a && \text{by induction} \\ &= \mathtt{fta}^{\rightarrowtail_1}_u(\alpha)) \hat{} \mathtt{I}_1(f)(\alpha, a) \\ &= \mathtt{fta}^{\rightarrowtail_1}_u(\alpha a). \end{aligned}$$

Case 3: $\mathtt{dom}_1(a) \not\rightarrowtail_1 u$ and $\mathtt{dom}_2(a) \not\rightarrowtail' r(u)$. In this case,

$$\begin{aligned} F_u(\mathtt{fta}^{\rightarrowtail'}_{r(u)}(\alpha a)) &= F_u(\mathtt{fta}^{\rightarrowtail'}_{r(u)}(\alpha)) \\ &= \mathtt{fta}^{\rightarrowtail_1}_u(\alpha) && \text{by induction} \\ &= \mathtt{fta}^{\rightarrowtail_1}_u(\alpha a). \end{aligned}$$

Case 4: $\mathtt{dom}_1(a) \not\rightarrowtail_1 u$ and $\mathtt{dom}_2(a) \overset{\top}{\rightarrowtail}' r(u)$. Then

$$\begin{aligned} F_u(\mathtt{fta}^{\rightarrowtail'}_{r(u)}(\alpha a)) &= F_u(\mathtt{fta}^{\rightarrowtail'}_{r(u)}(\alpha) \hat{} (\mathtt{fta}^{\rightarrowtail'}_{\mathtt{dom}_2(a)}(\alpha), a)) \\ &= F_u(\mathtt{fta}^{\rightarrowtail'}_{r(u)}(\alpha)) \\ &= \mathtt{fta}^{\rightarrowtail_1}_u(\alpha) \hat{} a && \text{by induction} \\ &= \mathtt{fta}^{\rightarrowtail_1}_u(\alpha a). \end{aligned}$$

□

Any machine M that is compliant with architectural specification $(\mathcal{DG}, \mathcal{C}_{\mathcal{DG}})$ does not reveal any C-dependent proposition to L. Since $\mathcal{C}_{\mathcal{DGR}}$ refines $\mathcal{C}_{\mathcal{DG}}$, the same property holds for any machine M that is compliant with $(\mathcal{DGR}, \mathcal{C}_{\mathcal{DGR}})$. We thus prove an information security property about architecture $\mathcal{DGR}$ by reference to the more abstract architecture $\mathcal{DG}$.

Theorem 16 *If M is FTA-compliant with architectural specification $(\mathcal{DGR}, \mathcal{C}_{\mathcal{DGR}})$ and $\pi(p)$ depends on C actions at α then $M, \pi, \alpha \vDash \neg K_L p$.*

Proof: Follows easily from Theorem 6 and Theorem 12, since $\mathcal{C}_{\mathcal{DGR}} \leq_r \mathcal{C}_{\mathcal{DG}}$. □

The proof of Theorem 15 uses the conditions T1-T3, and requires a somewhat laborious induction on the length of α to prove condition T3. The simpler and more local conditions for strict refinement cannot be used in this case, because $r(D) = D$, $r(L) = L$ and the edges $D \overset{\top}{\rightarrowtail} L$ in $\mathcal{DGR}$ and $D \overset{rel}{\rightarrowtail} L$ in $\mathcal{DGR}$ violate condition S2. To illustrate the application of strict refinement, we consider an architectural specification that varies $\mathcal{C}_{\mathcal{DGR}}$ by varying the allowed interpretations of the filter function f. According to $\mathcal{C}_{\mathcal{DGR}}$, *every* action of domain T is permitted to have an effect on domain H_P. In practice, some of the actions of domain T will have the purpose of communicating information to domain C, and it would not be desirable for such actions to be recorded in H_P. Let $\mathtt{ComTC}$ be the set of such communications actions. We specify that information about such actions is not permitted to flow from T to H_P. Let $\mathcal{C}'_{\mathcal{DGR}}$ be the architectural specification so that $(A, \mathtt{dom}, \mathtt{I}') \in \mathcal{C}'_{\mathcal{DGR}}$ if and only if $(A, \mathtt{dom}, \mathtt{I}')$ is an interpretation of $\mathcal{DGR}$ such that $\mathtt{I}'(f)$ is defined by $\mathtt{I}'(f)(\alpha, a) = a$ for all $\alpha \in A^*$ and $a \in A \setminus \mathtt{ComTC}$ with $\mathtt{dom}(a) = T$, and $\mathtt{I}'(f)(\alpha, a) = \epsilon$ otherwise. In particular, we now have, for the case where $\mathtt{dom}(a) = T$ and $a \in \mathtt{ComTC}$, that $\mathtt{I}'(f)(\alpha, a) = \epsilon$, whereas in the corresponding interpretation in $\mathcal{C}_{\mathcal{DGR}}$ we would have $\mathtt{I}(f)(\alpha, a) = a$ in this case.

Theorem 17 *$(\mathcal{DGR}, \mathcal{C}'_{\mathcal{DGR}})$ strictly refines $(\mathcal{DGR}, \mathcal{C}_{\mathcal{DGR}})$ by the identity function r.*[4]

Proof: Let $\mathcal{I}_1 = (A, \mathtt{dom}_1, \mathtt{I}_1) \in \mathcal{C}'_{\mathcal{DGR}}$. We may define an architectural interpretation for $\mathcal{DGR}$ by $\mathcal{I}_2 = (A, \mathtt{dom}_2, \mathtt{I}_2)$ where $\mathtt{dom}_2 = \mathtt{dom}_1$, and the interpretation $\mathtt{I}_2$ for $\mathcal{DGR}$ is defined by $\mathtt{I}_2(f)(\alpha, a) = a$ for all $\alpha \in A^*$ and $a \in A$ with $\mathtt{dom}_2(a) = T$, and $\mathtt{I}_2(f)(\alpha, a) = \epsilon$ otherwise.

Let $\mathcal{AI}_1 = (\mathcal{DGR}, \mathcal{I}_1)$ and $\mathcal{AI}_2 = (\mathcal{DGR}, \mathcal{I}_2)$. The domains of these two interpreted architectures are the same, so the identity function r on the set of domains $D_1 = D_2 = \{C, T, U, H_C, H_P, D, L\}$ has the right type to be a refinement mapping. We show that $\mathcal{AI}_1$ strictly refines $\mathcal{AI}_2$ by r, from which we may conclude that $(\mathcal{DGR}, \mathcal{C}'_{\mathcal{DGR}})$ strictly refines $(\mathcal{DGR}, \mathcal{C}_{\mathcal{DGR}})$ by r. For the proof, we use conditions S1-S3 and S5, and invoke Lemma 6 to obtain S4.

[4] We remark that this result implicitly also requires the technical side condition that $\epsilon \notin A$ for any set of actions A in an interpretation.

Conditions S1 and S2 are trivial from the fact that the architectures in $\mathcal{AI}_1$ and $\mathcal{AI}_2$ are the same and r is the identity function. For condition S3, suppose that $u, v \in D_1$ with $u \overset{f_1}{\rightarrowtail} v$ with $f_1 \neq \top$. Then $f_1 = f$, $u = T$ and $v = H_P$. We need to show that there exists an edge $r(u) \overset{f_2}{\rightarrowtail} r(v)$ with either $f_2 = \top$ or $f_2 \neq \top$ and for all $\alpha \in A^*$ and $a \in A$ with $\mathtt{dom}_1(a) = u$ we have $\mathtt{I}_2(f_2)(\alpha, a) = \epsilon$ implies $\mathtt{I}_1(f_1)(\alpha, a) = \epsilon$. Plainly, the edge $T \overset{f}{\rightarrowtail} H_P$ provides the required edge witness, with $f_2 = f \neq \top$, so it remains to check that for all $\alpha \in A^*$ and $a \in A$ with $\mathtt{dom}_1(a) = T$ we have $\mathtt{I}_2(f)(\alpha, a) = \epsilon$ implies $\mathtt{I}_1(f)(\alpha, a) = \epsilon$. This holds trivially, because $\mathtt{dom}_1(a) = T$ implies $\mathtt{dom}_2(a) = T$, in which case $\mathtt{I}_2(f)(\alpha, a) = a \neq \epsilon$.

For condition S5, let $a, b \in A$ and $v \in D_1$ such that $\mathtt{dom}_1(a) = \mathtt{dom}_1(b)$, and $\mathtt{dom}_1(a) \overset{f_1}{\rightarrowtail}_1 v$ and $\mathtt{dom}_2(a) \overset{f_2}{\rightarrowtail}_2 r(v)$, where $f_1 \neq \top$ and $f_2 \neq \top$. Then we must have $\mathtt{dom}_1(a) = \mathtt{dom}_2(b) = T$, $v = H_P$ and $f_1 = f_2 = f$. We need to show that for $\alpha, \beta \in A^*$, if $\mathtt{I}_2(f)(\alpha, a) = \mathtt{I}_2(f)(\beta, b) \neq \epsilon$, then $\mathtt{I}_1(f)(\alpha, a) = \mathtt{I}_1(f)(\beta, b)$. But $\mathtt{I}_2(f)(\alpha, a) = \mathtt{I}_2(f)(\beta, b) \neq \epsilon$ implies that $a = b \in A \setminus \mathtt{ComTC}$, so $\mathtt{I}_1(f)(\alpha, a) = a = b = \mathtt{I}_1(f)(\beta, b)$, as required.

To obtain condition S4 using Lemma 6, we need also that messages are source-identifying with respect to r. For this, define $\mathcal{S}(x) = \mathtt{dom}_1(x)$. Let $\alpha \in A^*$ and $a \in A$, and suppose $r(\mathtt{dom}_1(a)) \overset{f_2}{\rightarrowtail}_2 w$ and $\mathtt{I}_2(f_2)(\alpha, a) \neq \epsilon$. Then we have $f_2 = f$, $\mathtt{dom}_1(a) = T$ and $w = H_P$. Thus, $\mathtt{I}_2(f_2)(\alpha, a) = \mathtt{I}_2(f)(\alpha, a) = a$, and $\mathcal{S}(\mathtt{I}_2(f_2)(\alpha, a)) = \mathtt{dom}_1(a)$, as required to establish that messages are source-identifying. □

6 Implementing Architectures using Access Control

One of the mechanisms that might be used to enforce compliance with an information flow architecture is access control restrictions on the ability of domains to read and write objects. This idea was already implicit in the Bell and La Padula [1976] approach of enforcing that high level information should not flow to low level domains through a "no read up" and "no write down" access control policy. The idea was given a more semantically well-founded expression by Rushby [1992], who established a formal relation between access control systems and a theory of information flow based on intransitive noninterference policies. Rushby's "reference monitor conditions" give semantics to the notion of reading and writing, which was absent in the work of Bell and La Padula. Rushby's formulation was sharpened and shown to be closely related to TA-security by van der Meyden [2007].

In this section, we present a generalization of van der Meyden's formulation of access control, and show how enforcement of an access control policy together with local verification of trusted components can be used to assure that a system is compliant with an extended architecture.

We first recall some definitions and results from van der Meyden [2007]. The

system model we have used to this point does not require the states of a system to be equipped with any internal structure. In practice, systems typically will be constructed as an assembly of components. To capture this, Rushby [1992] introduced the notion of a *system with structured state*, which is a system M (with states S and domains D) together with

1. a set $\texttt{Obj}$ of *objects*,
2. a set V of *values*, and functions
3. $\texttt{contents} : S \times \texttt{Obj} \to V$, with $\texttt{contents}(s, n)$ interpreted as the value of object n in state s, and
4. $\texttt{observe}, \texttt{alter} : D \to \mathcal{P}(\texttt{Obj})$, with $\texttt{observe}(u)$ and $\texttt{alter}(u)$ interpreted as the set of objects that domain u can observe (or read) and alter (or write), respectively.

For brevity, we write $s(x)$ for $\texttt{contents}(s, x)$. We call the pair $(\texttt{observe}, \texttt{alter})$ the *access control* table of the machine. For each domain u, we define an equivalence relation of "observable content equivalence" on states $s, t \in S$ by $s \approx_u^{oc} t$ if $s(x) = t(x)$ for all $x \in \texttt{observe}(u)$.

Rushby introduced *reference monitor conditions* on such machines in order to capture formally the intuitions associated with the pair $(\texttt{observe}, \texttt{alter})$ being an access control table that restricts the ability of the actions to "read" and "write" the objects $\texttt{Obj}$. Van der Meyden [2007] sharpened these conditions to the following (the difference is in RM2):

RM1. If $s \approx_u^{oc} t$ then $\texttt{obs}_u(s) = \texttt{obs}_u(t)$.

RM2. For all actions $a \in A$, states $s, t \in S$ and objects $x \in \texttt{alter}(dom(a))$, if $s \approx_{\texttt{dom}(a)}^{oc} t$ and $s(x) = t(x)$ then $(s \cdot a)(x) = (t \cdot a)(x)$.

RM3. If $x \notin \texttt{alter}(dom(a))$ then $s(x) = (s \cdot a)(x)$

Intuitively, RM1 states that a domain's observation depends only on the values of the objects that it can observe (or read). RM2 states that if action a is performed in a domain u that is permitted to alter an object x, then the new value of the object after the action depends only on its old value and the values of objects that domain u is permitted to observe. The final conditions RM3 says that if action a is performed in a domain that is not permitted to alter (or write) an object x, then the value of x does not change.

We note that the terminology "reference monitor conditions" points to the fact that these conditions can be enforced by a reference monitor that mediates all attempts to perform an action, simply by denying requests by a domain u to read an object not in $\texttt{observe}(u)$ or write to an object not in $\texttt{alter}(u)$.

In addition to the reference monitor assumptions, Rushby considers a condition stating that if there is an object that may be altered by domain u and observed by domain v, then the information flow policy should permit flow of information from u to v. (Obviously, the object x provides a channel for information to flow from u to v.)

AOI. If $\texttt{alter}(u) \cap \texttt{observe}(v) \neq \emptyset$ then $u \rightarrowtail v$.

Van der Meyden [2007] shows the following, strengthening a result of Rushby [1992].

Theorem 18 *If M is a system with structured state satisfying RM1-RM3 and AOI with respect to $\rightarrowtail$ then M is TA-secure with respect to $\rightarrowtail$.*

We now develop a generalization of this result to extended architectures. As a first step, note that in extended architectures, the situation where the information flow policy potentially permits flow of information from domain u to domain v corresponds to the existence of an edge $u \rightarrowtail^{f} v$ for some label f (possibly $\top$). This motivates the following variant of condition AOI:

AOI′. If $\texttt{alter}(u) \cap \texttt{observe}(v) \neq \emptyset$ then $u \rightarrowtail^{f} v$ for some f.

Next, we develop a set of conditions that check that information flow constraints of the form $u \rightarrowtail^{f} v$ with $f \neq \top$ have been correctly implemented in a system. Let $\mathcal{I} = (A, \texttt{dom}, \texttt{I})$ be an interpretation of architecture $\mathcal{A} = (D, \rightarrowtail)$. Consider the following constraints in a system M with actions A, domains D and domain function $\texttt{dom}$:

I1. If $\texttt{dom}(a) \rightarrowtail^{f} u$ for $f \neq \top$ and $\texttt{I}(f)(\alpha, a) = \epsilon$ and $x \in \texttt{observe}(u) \cap \texttt{alter}(\texttt{dom}(a))$ then $(s_0 \cdot \alpha a)(x) = (s_0 \cdot \alpha)(x)$.

I2. If $\texttt{dom}(a) \rightarrowtail^{f} u$ with $f \neq \top$ and $\texttt{dom}(b) \rightarrowtail^{g} u$ with $f \neq \top$ and $\texttt{I}(f)(\alpha, a) = \texttt{I}(g)(\beta, b) \neq \epsilon$ and $x \in \texttt{observe}(u) \cap (\texttt{alter}(\texttt{dom}(a)) \cup \texttt{alter}(\texttt{dom}(b)))$ and $(s_0 \cdot \alpha)(x) = (s_0 \cdot \beta)(x)$ then $(s_0 \cdot \alpha a)(x) = (s_0 \cdot \beta b)(x)$.

Condition I1 ensures that if filter function f restricts how the domain of action a may interact with domain u (i.e., $\texttt{dom}(a) \rightarrowtail^{f} u$ with $f \neq \top$), and the filter function interpretation does not allow any information flow ($\texttt{I}(f)(\alpha, a) = \epsilon$), then the action does not change the state of any object x that domain u may observe and domain $\texttt{dom}(a)$ is allowed to alter. Condition I2 states that if an action a in domain $\texttt{dom}(a)$ may alter an object x that is observable by domain u and $\texttt{dom}(a) \rightarrowtail^{f} u$ then the new state of object x is determined by the interpretation of filter function f.

We note that verification of these constraints requires consideration only of domains that are trusted, in the sense that they have outgoing edges not labelled $\top$, and the objects that such domains are permitted to alter. Thus verification of these constraints can be localized to the trusted domains. The following result states that such local verification, together with enforcement of an access control policy consistent with the information flow policy via a mechanism satisfying the reference monitor constraints, suffices to assure that an information flow policy has been satisfied:

Theorem 19 *Let $\mathcal{AI}$ be a strongly non-conflating interpreted architecture. Suppose that M is a system with structured state satisfying RM1-RM3, AOI′ and I1-I2. Then M is FTA-compliant with $\mathcal{AI}$.*

Proof: We show that $\mathtt{fta}_u(\alpha) = \mathtt{fta}_u(\beta)$ implies that $s_0 \cdot \alpha \approx_u^{oc} s_0 \cdot \beta$, for all domains u and $\alpha, \beta \in A^*$. Note that it then follows using RM1 that $\mathtt{fta}_u(\alpha) = \mathtt{fta}_u(\beta)$ implies $\mathtt{obs}_u(\alpha) = \mathtt{obs}_u(\beta)$, which shows that M is FTA-compliant with $\mathcal{AI}$.

The proof proceeds by induction on the combined length of α and β. The base case of $\alpha = \beta = \epsilon$ is trivial. Suppose that $\mathtt{fta}_u(\alpha a) = \mathtt{fta}_u(\beta)$, where the statement holds for sequences of shorter combined length. We consider several cases:

Case 1: $\mathtt{dom}(a) \not\rightarrow u$. In this case $\mathtt{fta}_u(\alpha) = \mathtt{fta}_u(\alpha a) = \mathtt{fta}_u(\beta)$, so by induction we have $s_0 \cdot \alpha \approx_u^{oc} s_0 \cdot \beta$. We need to show that $s_0 \cdot \alpha a \approx_u^{oc} s_0 \cdot \beta$, i.e., that $(s_0 \cdot \alpha a)(x) = (s_0 \cdot \beta)(x)$ for all $x \in \mathtt{observe}(u)$. Let $x \in \mathtt{observe}(u)$. Since $\mathtt{dom}(a) \not\rightarrow u$, it follows using AOI′ that $x \notin \mathtt{alter}(\mathtt{dom}(a))$. Thus, by RM3, we obtain that $(s_0 \cdot \alpha a)(x) = ((s_0 \cdot \alpha) \cdot a)(x) = (s_0 \cdot \alpha)(x) = (s_0 \cdot \beta)(x)$, as required.

Case 2: $\mathtt{dom}(a) \xrightarrow{\top} u$. In this case $\mathtt{fta}_u(\beta) = \mathtt{fta}_u(\alpha a) = \mathtt{fta}_u(\alpha)\,\hat{}\,(\mathtt{fta}_{\mathtt{dom}(a)}(\alpha), a)$, so $\beta \neq \epsilon$. Let $\beta = \gamma b$. If $\mathtt{dom}(b) \not\rightarrow u$, then we can swap the roles of αa and γb and apply Case 1. If $\mathtt{dom}(b) \xrightarrow{f} u$ with $f \neq \top$, then $\mathtt{fta}_u(\gamma b) = \mathtt{fta}_u(\gamma)\,\hat{}\,\mathtt{I}(f)(\gamma, b)$, and we obtain that $\mathtt{I}(f)(\gamma, b) = (\mathtt{fta}_{\mathtt{dom}(a)}(\alpha), a)$. Since the architecture is strongly non-conflating, this case is not possible. We are left with the case that $\mathtt{dom}(b) \xrightarrow{\top} u$. Here $\mathtt{fta}_u(\gamma b) = \mathtt{fta}_u(\gamma)\,\hat{}\,(\mathtt{fta}_{\mathtt{dom}(b)}(\gamma), b)$, and we conclude that $\mathtt{fta}_u(\gamma) = \mathtt{fta}_u(\gamma)$ and $\mathtt{fta}_{\mathtt{dom}(a)}(\alpha) = \mathtt{fta}_{\mathtt{dom}(b)}(\gamma)$ and $a = b$. By the induction hypothesis, we obtain that $s_0 \cdot \alpha \approx_u^{oc} s_0 \cdot \gamma$ and $s_0 \cdot \alpha \approx_{\mathtt{dom}(a)}^{oc} s_0 \cdot \gamma$. The former states that for all objects $x \in \mathtt{observe}(u)$, we have $(s_0 \cdot \alpha)(x) = (s_0 \cdot \gamma)(x)$. Thus, by RM2 and the fact that $a = b$, it follows that $(s_0 \cdot \alpha a)(x) = ((s_0 \cdot \alpha) \cdot a)(x) = ((s_0 \cdot \gamma) \cdot b)(x) = (s_0 \cdot \gamma b)(x)$ for all $x \in \mathtt{observe}(u)$. This shows that $(s_0 \cdot \alpha a) \approx_u^{oc} (s_0 \cdot \gamma b)$, as required.

Case 3: $\mathtt{dom}(a) \xrightarrow{f} u$. In this case we have

$$\mathtt{fta}_u(\beta) = \mathtt{fta}_u(\alpha a) = \mathtt{fta}_u(\alpha)\,\hat{}\,\mathtt{I}(f)(\alpha, a).$$

There are several possibilities:

- Case 3a: $\mathtt{I}(f)(\alpha, a) = \epsilon$. In this case we obtain $\mathtt{fta}_u(\alpha) = \mathtt{fta}_u(\beta)$, so $s_0 \cdot \alpha \approx_u^{oc} s_0 \cdot \beta$ by induction. By RM3, we have $(s_0 \cdot \alpha a)(x) = ((s_0 \cdot \alpha) \cdot a)(x) = (s_0 \cdot \alpha)(x)$ for $x \in \mathtt{observe}(u)$ such that $x \notin \mathtt{alter}(\mathtt{dom}(a))$. Further, by I1, we have $(s_0 \cdot \alpha a)(x) = (s_0 \cdot \alpha)(x)$ for $x \in \mathtt{observe}(u) \cap \mathtt{alter}(\mathtt{dom}(a))$. Thus, $s_0 \cdot \alpha a \approx_u^{oc} s_0 \cdot \alpha$, and we conclude that $s_0 \cdot \alpha a \approx_u^{oc} s_0 \cdot \beta$, as required.

- Case 3b: $\mathtt{I}(f)(\alpha, a) \neq \epsilon$. Here we must have $\beta \neq \epsilon$, so let $\beta = \gamma b$. Arguing as above, we cannot have $\mathtt{dom}(b) \xrightarrow{\top} u$ since that contradicts the assumption that the architecture is strongly non-conflating. The case that $\mathtt{dom}(b) \not\rightarrow u$ can be handled using Case 1 above. We are left with the possibility that $\mathtt{dom}(b) \xrightarrow{g} u$ for some $g \neq \top$. Here $\mathtt{fta}_u(\gamma b) = \mathtt{fta}_u(\gamma)\,\hat{}\,\mathtt{I}(g)(\gamma, b)$. The

situation where $\mathtt{I}(g)(\gamma, b) = \epsilon$ can be handled using the argument of Case 3a. Thus, we may assume that $\mathtt{I}(g)(\gamma, b) \neq \epsilon$. We then obtain that $\mathtt{fta}_u(\alpha) = \mathtt{fta}_u(\gamma)$ and $\mathtt{I}(f)(\alpha, a) = \mathtt{I}(g)(\gamma, b) \neq \epsilon$. By the induction hypothesis, we have $s_0 \cdot \alpha \approx_u^{oc} s_0 \cdot \gamma$, so $(s_0 \cdot \alpha)(x) = (s_0 \cdot \gamma)(x)$ for all $x \in \mathtt{observe}(u)$. We need to show that $(s_0 \cdot \alpha a)(x) = (s_0 \cdot \gamma b)(x)$ for all $x \in \mathtt{observe}(u)$. If $x \in \mathtt{observe}(u) \cap (\mathtt{alter}(\mathtt{dom}(a) \cup \mathtt{alter}(\mathtt{dom}(b))$, then this follows using I2. On the other hand, if $x \notin \mathtt{alter}(\mathtt{dom}(a))$ and $x \notin \mathtt{alter}(\mathtt{dom}(b))$, then by RM3 we get $(s_0 \cdot \alpha a)(x) = ((s_0 \cdot \alpha) \cdot a)(x) = (s_0 \cdot \alpha)(x) = (s_0 \cdot \gamma)(x) = ((s_0 \cdot \gamma) \cdot b)(x) = (s_0 \cdot \gamma b)(x)$, so the desired conclusion holds in either case.

□

Conditions I1, I2 are still somewhat non-local because of the reference to α. Since architectural specifications are stated in terms of these state sequences, references to them cannot be completely eliminated. However, it is often convenient to factor this reference via properties of the state of the system. Suppose that for the edge labels f there exist functions $F(f)$ with domain $S \times A$, satisfying the following constraints:

E1. $\mathtt{I}(f)(\alpha, a) = \mathtt{I}(g)(\beta, b)$ implies $F(f)(s_0 \cdot \alpha, a) = F(g)(s_0 \cdot \beta, b)$

E2. $\mathtt{I}(f)(\alpha, a) = \epsilon$ implies $F(f)(s_0 \cdot \alpha, a) = \epsilon$.

Intuitively, F gives a state-based encoding of an approximation to $\mathtt{I}$.

The following conditions use F to give a variant of the conditions I1, I2 that is stated with respect to the states of the machine:

I1′. If $\mathtt{dom}(a) \rightarrowtail^{f} u$ for $f \neq \top$ and $F(f)(s, a) = \epsilon$ and $x \in \mathtt{observe}(u) \cap \mathtt{alter}(\mathtt{dom}(a))$ then $(s \cdot a)(x) = (s)(x)$.

I2′. If $\mathtt{dom}(a) \rightarrowtail^{f} u$ with $f \neq \top$ $\mathtt{dom}(b) \rightarrowtail^{g} u$ with $g \neq \top$ and $F(f)(s, a) = F(g)(t, b) \neq \epsilon$ and $x \in \mathtt{observe}(u) \cap (\mathtt{alter}(\mathtt{dom}(a)) \cup \mathtt{alter}(\mathtt{dom}(b)))$ and $s(x) = t(x)$ then $(s \cdot a)(x) = (t \cdot b)(x)$.

The following result states that this state-factored representation implies the conditions I1-I2.

Theorem 20 *Conditions E1-E1 and I1′-I2′ imply I1 and I2.*

Proof: Assume E1-E1 and I1′-I2′. For I1, suppose that $\mathtt{dom}(a) \rightarrowtail^{f} u$ for $f \neq \top$ and $\mathtt{I}(f)(\alpha, a) = \epsilon$ and let $x \in \mathtt{observe}(u) \cap \mathtt{alter}(\mathtt{dom}(a))$. We need to show that $(s_0 \cdot \alpha a)(x) = (s_0 \cdot \alpha)(x)$. From $\mathtt{I}(f)(\alpha, a) = \epsilon$, we get by E2 that $F(s_0 \cdot \alpha, a) = \epsilon$. By I1′ with $s = s_0 \cdot \alpha$, we the conclude that $(s_0 \cdot \alpha a)(x) = (s_0 \cdot \alpha)(x)$, as required.

For I2, suppose that $\mathtt{dom}(a) \rightarrowtail^{f} u$ with $f \neq \top$ and $\mathtt{dom}(b) \rightarrowtail^{g} u$ with $f \neq \top$ and $\mathtt{I}(f)(\alpha, a) = \mathtt{I}(g)(\beta, b) \neq \epsilon$ and $x \in \mathtt{observe}(u) \cap (\mathtt{alter}(\mathtt{dom}(a)) \cup \mathtt{alter}(\mathtt{dom}(b)))$ and $(s_0 \cdot \alpha)(x) = (s_0 \cdot \beta)(x)$. We need to show that $(s_0 \cdot \alpha a)(x) = (s_0 \cdot \beta b)(x)$. By E1, we have $F(\alpha, a) = F(\beta, b)$. If both values are equal to ϵ, then

we use I1′ to conclude that $(s_0 \cdot \alpha a)(x) = (s_0 \cdot \alpha)(x)$ and $(s_0 \cdot \beta b)(x) = (s_0 \cdot \beta)(x)$. It then follows that $(s_0 \cdot \alpha a)(x) = (s_0 \cdot \beta b)(x)$, as required. Alternately, if $F(\alpha, a) = F(\beta, b) \neq \epsilon$, then by I2′, with $s = s_0 \cdot \alpha$ and $t = s_0 \cdot \beta$, we conclude that $(s_0 \cdot \alpha a)(x) = (s \cdot a)(x) = (t \cdot b)(x) = (s_0 \cdot \beta b)(x)$, as required. □

To illustrate the application of these results, we consider the examples of architectural specifications introduced above, and show how some particular systems may be proved to be implementations of these specifications.

6.1 Starlight Architecture

We present an implementation of the Starlight architectural specification $(\mathcal{SL}, \mathcal{C}_{\mathcal{SL}})$ as a system with structured state. We first select an architectural interpretation $\mathcal{I} = (A, \mathtt{dom}, \mathtt{I})$ in $\mathcal{C}_{\mathcal{SL}}$. The actions A of the interpretation are given with their associated domain as follows:

- The actions of domain S consist of the toggle action $\mathtt{t}$, plus actions k from a set $\mathcal{K}$, which intuitively represents the set of keyboard actions.
- The actions of domain H consist of the action $\mathtt{get}$, plus actions h drawn from some set E_H representing possible High level events.
- The actions of domain L consist of the action $\mathtt{put}$, plus actions l drawn from some set E_L representing the possible Low level events.

As required by $\mathcal{C}_{\mathcal{SL}}$, we assume that $\mathtt{I}$ satisfies

$$\mathtt{I}(sf)(\alpha, a) = \begin{cases} a & \text{if } a = \mathtt{t} \text{ or} \\ & \quad \#_{\mathtt{t}}(\alpha) \text{ is odd and } a \in A_S \\ \epsilon & \text{otherwise} \end{cases}$$

where $\#_{\mathtt{t}}(\alpha)$ is the number of occurrences of t in α.

Next, we construct a system M for the interpreted architecture $(\mathcal{SL}, \mathcal{I})$. The system M is based on the set of objects $\mathtt{Obj} = \{logH, logL, togpos\}$. Intuitively, $logH$ and $logL$ are logs of actions observable to H and L, respectively, and $togpos$ indicates the position of the toggle switch. The objects $logH$ and $logL$ have values in A^*, and $togpos$ has a value in $\{H, L\}$. A state s is just an assignment of value of the given type to each of the objects, and we have $\mathtt{contents}(s, x) = s(x)$. The initial state s_0 is defined by $s(logH) = s(logL) = \epsilon$ and $s(togpos) = H$. Transitions are defined by the following code associated to each of the actions:

- for domain S:

  $\mathtt{t}$: if $togpos = H$ then $togpos := L$ else $togpos := H$

  k: if $togpos = H$ then $logH := logH \hat{} k$ else $logL := logL \hat{} k$

  That is, the toggle action $\mathtt{t}$ changes the position of the toggle, and k is recorded in either the High level log or the Low level log, depending on the position of the toggle.

- for domain H:

  h: $logH := logH \hat{} h$

  `get`: $logH := logH \hat{} logL$

  Thus, the events h are recorded in the High level log, and `get` fetches a copy of the low level log.

- for domain L:

  l: $logL := logL \hat{} l$

  `put`: $logH := logH \hat{} logL$

  Similarly the events l are recorded in the Low level log, and `put` pushes a copy of the low level log into the High level log.

The access control functions `alter` and `observe` are defined by the following access control table. Here, for an object x and a domain u, we have an entry a

	H	L	S
$logH$	a,o	a	a
$logL$	o	a,o	a
$togpos$			a,o

(or o) in the corresponding cell of the table just when $x \in \texttt{alter}(u)$ (respectively, $x \in \texttt{observe}(u)$). We define observations in the system by allowing each domain to observe the values of all of its observable objects. That is, for each domain u and state s, we define $\texttt{obs}_u(s) = s \restriction \texttt{observe}(u)$.

Proposition 2 *The system M is FTA-compliant with architectural specification $(\mathcal{SL}, \mathcal{C}_{\mathcal{SL}})$.*

Proof: We show that M satisfies RM1-RM3, AOI′ and I1-I2 with respect to $(\mathcal{SL}, \mathcal{I})$, and invoke Theorem 19 to conclude FTA-compliance. RM1 is immediate from the fact that we have defined $\texttt{obs}_u(s) = s \restriction \texttt{observe}(u)$.

For RM2, we consider each of the actions a and $x \in \texttt{alter}(\texttt{dom}(a))$ in turn, and show that if $s \approx^{oc}_{\texttt{dom}(a)} t$ and $s(x) = t(x)$ then $(s \cdot a)(x) = (t \cdot a)(x)$. Note that we need to consider only those $x \in \texttt{alter}(\texttt{dom}(a))$ whose value could be changed by action a, since otherwise $(s \cdot a)(x) = (t \cdot a)(x)$ is immediate from $s(x) = t(x)$. (Since the only way an action could change an object, given the type of code we have used to define the transition relation is by use of the assignment operator, for each action we can take the set of objects that it could change to be the set of objects to which the code for the action makes an assignment.)

- *Actions a with $\texttt{dom}(a) = S$.* Suppose $s \approx^{oc}_{S} t$. Since $\texttt{observe}(S) = \{togpos\}$, this means $s(togpos) = t(togpos)$. Because $\texttt{alter}(S) = \texttt{Obj}$, we need to consider all objects $x \in \texttt{Obj}$ that could be changed by a, assuming that $s(x) = t(x)$.

The only object in $\texttt{alter}(S)$ that could be changed by action $\texttt{t}$ is *togpos*. If $s(togpos) = t(togpos) = H$, we have $(s \cdot \texttt{t})(togpos) = L = (t \cdot \texttt{t})(togpos)$, as required, and similarly if $s(togpos) = t(togpos) = L$.

For action k, the only objects x that could be changed by k are *logH* and *logL*. We consider the case of *logL*. Assume $s(logL) = t(logL)$. If $s(togpos) = t(togpos) = H$ then $(s \cdot k)(logL) = s(logL) = t(logL) = (t \cdot k)(logL)$, as required. If $s(togpos) = t(togpos) = L$, then $(s \cdot k)(logL) = s(logL)\hat{}\, k = t(logL)\hat{}\, k = (t \cdot k)(logL)$. The argument for *logH* is similar.

- *Actions a with* $\texttt{dom}(a) = H$. Assume $s \approx^{oc}_{H} t$. By the access control table, this means that $s(logH) = t(logH)$ and $s(logL) = t(logL)$. For actions h and $\texttt{get}$, note that the only object these actions could change is *logH*. We have $(s \cdot h)(logH) = s(logH)\hat{}\, k = t(logH)\hat{}\, k = (t \cdot h)(logH)$. Similarly, $(s \cdot h)(logH) = s(logH)\hat{}\, s(logL) = t(logH)\hat{}\, t(logL) = (t \cdot h)(logH)$.

- *Actions a with* $\texttt{dom}(a) = L$. Assume $s \approx^{oc}_{L} t$. By the access control table, this means that $s(logL) = t(logL)$.

  Action l could change only *logL*. We have $(s \cdot l)(logL) = s(logL)\hat{}\, l = t(logL)\hat{}\, l = (t \cdot l)(logL)$.

  Action $\texttt{put}$ could change only *logH*. Assume that $s(logH) = t(logH)$. Then $(s{\cdot}\texttt{put})(logH) = s(logH)\hat{}\, s(logL) = t(logH)\hat{}\, t(logL) = (t{\cdot}\texttt{put})(logH)$.

This completes the proof that RM2 is satisfied.

For RM3, we need to show that if $x \notin \texttt{alter}(\texttt{dom}(a))$, then $(s \cdot a)(x) = s(x)$. This can be done by inspection of the code for each action. In the case of $\texttt{dom}(a) = S$, we have $\texttt{Obj} \setminus \texttt{alter}(\texttt{dom}(a)) = \emptyset$, so the claim is trivial. For $\texttt{dom}(a) = H$, we have $\texttt{Obj} \setminus \texttt{alter}(\texttt{dom}(a)) = \{logL, togpos\}$, but the only object assigned to by H actions is *logH*, so the values of objects in this set are not changed by H actions. For $\texttt{dom}(a) = L$, we have $\texttt{Obj} \setminus \texttt{alter}(\texttt{dom}(a)) = \{togpos\}$, and L actions do not assign to *togpos*, so again there is no change to the value of this object.

To show AOI′, we consider the possible values for $x \in \texttt{alter}(u) \cap \texttt{observe}(v)$, and check that there exists an edge $u \xrightarrow{f} v$ in each case.

- For $x = logH$, the only case of $x \in \texttt{observe}(v)$ is $v = H$. Since there is an edge from every domain v to H (in the case of $v = H$ there is the implicit reflexive edge labelled $\top$), AOI′ is immediate in this case.

- For $x = logL$, we have that $x \in \texttt{alter}(u)$ just when $u \in \{L, S\}$, and $x \in \texttt{observe}(v)$ just when $v \in \{L, H\}$. For each of these possible values of u, v there exists an edge from u to v.

- for $x = togpos$, we have $x \in \texttt{alter}(u) \cap \texttt{observe}(v)$ just when $u = v = S$, in which case we have the reflexive edge.

Before considering I1-I2, we establish the following invariant on the states of M:

$$\text{For all } \alpha \in A^*, \text{ the number } \#_{\texttt{t}}(\alpha) \text{ is even iff } (s_0 \cdot \alpha)(togpos) = H. \qquad (1)$$

It is worth remarking that the proof of this invariant itself benefits from the property RM3 that we have just established. Note that $\mathit{togpos} \in \mathtt{alter}(u)$ iff $u = S$. Thus, using RM3, we have that $(s_0 \cdot \alpha)(\mathit{togpos}) = s_0 \cdot (\alpha \restriction \{S\})(\mathit{togpos})$. Moreover, inspection of the code shows that the actions k of domain S do not change the value of togpos. Thus, in fact $(s_0 \cdot \alpha)(\mathit{togpos}) = s_0 \cdot (\alpha \restriction \{\mathtt{t}\})(\mathit{togpos})$, where we interpret $\alpha \restriction B$ where B is a set of actions to mean the subsequence of α of actions in B. A straightforward induction on sequences $\beta \in \{\mathtt{t}\}^*$ shows that $(s_0 \cdot \beta)(\mathit{togpos}) = H$ if $\#_{\mathtt{t}}(\beta)$ is even and $(s_0 \cdot \beta)(\mathit{togpos}) = L$ otherwise. Since $\#_{\mathtt{t}}(\alpha) = \#_{\mathtt{t}}(\alpha \restriction \{\mathtt{t}\})$, we have $\#_{\mathtt{t}}(\alpha)$ even iff $\#_{\mathtt{t}}(\alpha \restriction \{\mathtt{t}\})$ is even iff $(s_0 \cdot \alpha \restriction \{\mathtt{t}\})(\mathit{togpos}) = H$ iff $(s_0 \cdot \alpha)(\mathit{togpos}) = H$, as required.

We now prove I1. The only case where $\mathtt{dom}(a) \rightarrowtail^{f} u$ for $f \neq \top$ is $\mathtt{dom}(a) = S$ and $u = L$ and $f = \mathit{sf}$. Suppose that $\mathtt{I}(\mathit{sf})(\alpha, a) = \epsilon$ and $x \in \mathtt{observe}(L) \cap \mathtt{alter}(S)$. We need to show that $(s_0 \cdot \alpha a)(x) = (s_0 \cdot \alpha)(x)$. Now $\mathtt{observe}(L) \cap \mathtt{alter}(S) = \{\mathit{logL}\}$, so we have $x = \mathit{logL}$. Further, if $\mathtt{dom}(a) = S$ then $\mathtt{I}(\mathit{sf})(\alpha, a) = \epsilon$ just when $a \neq \mathtt{t}$ and $\#_{\mathtt{t}}(\alpha)$ is even. It therefore suffices to consider the case of $a = k$. By the invariant (1), since $\#_{\mathtt{t}}(\alpha)$ is even, we have $(s_0 \cdot \alpha)(\mathit{togpos}) = H$, so the code for $a = k$ takes the else branch, which does not change $x = \mathit{logL}$. Hence $(s_0 \cdot \alpha a)(x) = (s_0 \cdot \alpha)(x)$, as required.

For I2, the only case where $\mathtt{dom}(a) \rightarrowtail^{f} u$ with $f \neq \top$ and $\mathtt{dom}(b) \rightarrowtail^{g} u$ with $g \neq \top$ is $\mathtt{dom}(a) = \mathtt{dom}(b) = S$ and $u = L$ and $f = g = \mathit{sf}$. Suppose that $\mathtt{I}(f)(\alpha, a) = \mathtt{I}(g)(\beta, b) \neq \epsilon$ and $x \in \mathtt{observe}(u) \cap (\mathtt{alter}(\mathtt{dom}(a)) \cup \mathtt{alter}(\mathtt{dom}(b)))$ and $(s_0 \cdot \alpha)(x) = (s_0 \cdot \beta)(x)$. We need to show that $(s_0 \cdot \alpha a)(x) = (s_0 \cdot \beta b)(x)$. In the relevant case, $\mathtt{observe}(u) \cap (\mathtt{alter}(\mathtt{dom}(a)) \cup \mathtt{alter}(\mathtt{dom}(b))) = \mathtt{observe}(L) \cap \mathtt{alter}(S) = \{\mathit{logL}\}$, so we have $x = \mathit{logL}$. Now, if $\mathtt{I}(f)(\alpha, a) = \mathtt{I}(g)(\beta, b) \neq \epsilon$ then $a = b = \mathtt{t}$ or both $\#_{\mathtt{t}}(\alpha)$ and $\#_{\mathtt{t}}(\beta)$ are odd and $a = b \in A_S \setminus \{\mathtt{t}\}$. In the former case, the code for $a = b = \mathtt{t}$ does not change the value of $x = \mathit{logL}$, so it is immediate from $(s_0 \cdot \alpha)(x) = (s_0 \cdot \beta)(x)$ that $(s_0 \cdot \alpha a)(x) = (s_0 \cdot \beta b)(x)$. In the latter case, we have by the invariant (1) that $(s_0 \cdot \alpha)(\mathit{togpos}) = (s_0 \cdot \beta)(\mathit{togpos}) = L$, so the code for $a = k$ takes the else branch from both states and we have $(s_0 \cdot \alpha a)(\mathit{logL}) = (s_0 \cdot \alpha)(\mathit{logL}) \hat{} a = (s_0 \cdot \beta)(\mathit{logL}) \hat{} a = (s_0 \cdot \beta a)(\mathit{logL})$, as required. □

Intuitively, the proof of RM2 can be summarized to be a consequence of the fact that for each action a, the code for action a "reads" only objects in $\mathtt{observe}(\mathtt{dom}(a))$, and "writes" to objects in $\mathtt{alter}(\mathtt{dom}(a))$ in ways that depend only on the values read and the values of the object written to. Note that the action `put` assigns a value to the H object logH that depends on the value of the same object logH, which L is not permitted to observe/read. However, this is not a violation of RM2. We expect that it is possible to develop a static analysis that would be sound for checking RM1-3, but we do not pursue this here.

6.2 Downgrader Architecture

Next, we give an example that illustrates the role that code structure can play in the enforcement of an architecture. We sketch a class of implementations of the refined downgrader architectural specification $(\mathcal{DGR}, \mathcal{C}_{\mathcal{DGR}})$ of Section 5.4,

in the form of systems with structured state. Let $\mathcal{I} = (A, \mathtt{dom}, \mathtt{I})$ be an architectural interpretation in $\mathcal{C}_{\mathcal{DGR}}$. The set of actions A_u for each domain u is some arbitrary set, and $\mathtt{I}(f)$ is defined by $\mathtt{I}(f)(\alpha, a) = a$ for all $\alpha \in A^*$ and $a \in A$ with $\mathtt{dom}(a) = T$, and $\mathtt{I}(f)(\alpha, a) = \epsilon$ otherwise.

We describe the system M abstractly: for each domain u we assume that some set of objects $\mathtt{Obj}_u$ is given. For the domain H_P, we take $\mathtt{Obj}_{H_P} = \{\mathit{buff}, \mathit{logP}\}$. Intuitively, logP is a log of events that are observable to the domain H_P, and buff is a buffer storing declassification requests from L. The other domains C, T, U, H_C, D and L are associated with an arbitrary set of objects. The set of all objects $\mathtt{Obj}$ is the union of these sets $\mathtt{Obj}_u$.

The access control table on these objects is given in the following "capability" formatted table.

	C	T	U	H_C	H_P	D	L
observe	$\mathtt{Obj}_C$	$\mathtt{Obj}_T$	$\mathtt{Obj}_U$	$\mathtt{Obj}_{H_C}$	logP	$\mathit{buff}, \mathtt{Obj}_D$	$\mathtt{Obj}_L$
alter	$\mathtt{Obj}_C$ $\mathtt{Obj}_{H_C}$	$\mathtt{Obj}_T$ $\mathtt{Obj}_{H_C}$ logP	$\mathtt{Obj}_U$ logP	$\mathtt{Obj}_{H_C}$ $\mathtt{Obj}_T$ $\mathtt{Obj}_C$	$\mathtt{Obj}_T$,$\mathtt{Obj}_U$ $\mathtt{Obj}_D, \mathit{logP}$	$\mathtt{Obj}_D$,$\mathtt{Obj}_L$ $\mathit{buff}, \mathit{logP}$	$\mathtt{Obj}_L$ buff

For example, this table says that an object $x \in \mathtt{Obj}_U$ can be observed by domain U only, and can be altered by domains U and H_P.

States are defined to be assignments of a value in some set for each $x \in \mathtt{Obj}$. In case of buff, the value $s(\mathit{buff})$ is assumed to be in the set A_L^* of L action sequences. In case of $\log_P$, the value is assumed to be in the set $(A_T \cup A_U \cup A_{H_P})^*$, i.e., this log records actions in the domains T, U and H_P. We assume each domain observes all its observable objects, i.e., $\mathtt{obs}_u(s) = s \restriction \mathtt{observe}(u)$.

In order to focus on the semantic conditions associated to the edge $T \overset{f}{\rightarrowtail} H_P$, we assume that the actions for all domains T are given semantics in such a way as to satisfy the reference monitor conditions RM2-RM3. For actions $a \in A_T$, we assume that the code for these actions (which determines the state transition when the action is performed) has the following structure:

$$\mathit{logP} := \mathit{logP} \,\hat{}\, a; P_a$$

where P_a is a program that reads only objects in $\mathtt{Obj}_T$ and writes only objects in $\mathtt{Obj}_T \cup \mathtt{Obj}_{H_C}$.

We now argue that a system within this class of systems satisfies conditions RM1-RM3, AOI′ and I1-I2. Condition RM1 is trivial from the definition of observations. Conditions RM2-RM3 have been assumed for all actions except $a \in A_T$. We argue that the latter also satisfy RM2. Given the allowed code pattern for such actions a, the only objects x that are altered by a are in $\{\mathit{logP}\} \cup \mathtt{Obj}_T \cup \mathtt{Obj}_{H_C}$. The object logP is altered by appending a. This gives a result that depends only on the previous value of logP (and the action a being performed), which satisfies RM2. Other objects $x \in \mathtt{Obj}_T \cup \mathtt{Obj}_{H_C}$ are altered by the code in ways that depend only on the action a and the value of objects in $\mathtt{Obj}_T = \mathtt{observe}(T)$, which again satisfies RM2.

For RM3 in the case of actions $a \in A_T$, note that the complement of $\mathtt{alter}(T)$ is $\mathtt{Obj}_C \cup \mathtt{Obj}_U \cup \mathtt{Obj}_D \cup \mathtt{Obj}_L \cup \{\mathit{buff}\}$. Since none of these objects are written by the code pattern for a, RM3 is satisfied in this case.

Satisfaction of AOI′ can be checked by inspection, considering all possibilities for $x \in \mathtt{alter}(u) \cap \mathtt{observe}(v)$ in turn and verifying that there is an edge from u to v in the policy in each case. For example, we have $\mathit{logP} \in \mathtt{alter}(T) \cap \mathtt{observe}(H_C)$, but we also have the edge $T \rightarrowtail^{f} H_C$.

To check conditions I1-I2, we need only consider the edge $T \rightarrowtail^{f} H_C$, i.e., we consider actions a with $\mathtt{dom}(a) = T$ and let $u = H_C$. Condition I1 has precondition $\mathtt{I}(f)(\alpha, a) = \epsilon$. Since always $\mathtt{I}(f)(\alpha, a) = a \neq \epsilon$, the precondition for I1 is always false and this condition is trivially true.

For condition I2, suppose that $\mathtt{dom}(a) = \mathtt{dom}(b) = T$ and $\mathtt{I}(f)(\alpha, a) = \mathtt{I}(f)(\beta, b) \neq \epsilon$ and $x \in \mathtt{observe}(H_P) \cap \mathtt{alter}(T)$ and $(s_0 \cdot \alpha)(x) = (s_0 \cdot \beta)(x)$. By definition of $\mathtt{I}(f)$ we have $a = b$. Since $\mathtt{observe}(H_P) \cap \mathtt{alter}(T) = \{\mathit{logP}\}$, we have $x = \mathit{logP}$ and $(s_0 \cdot \alpha)(\mathit{logP}) = (s_0 \cdot \beta)(\mathit{logP})$. Since the code $P_a = P_b$ does not alter logP, we obtain that $(s_0 \cdot \alpha a)(\mathit{logP}) = (s_0 \cdot \alpha)(\mathit{logP}) \hat{} a = (s_0 \cdot \beta)(\mathit{logP}) \hat{} b = (s_0 \cdot \beta b)(\mathit{logP})$, as required.

This completes the argument that a system constructed as described FTA-complies with the refined downgrader architectural specification $(\mathcal{DGR}, \mathcal{C}_{\mathcal{DGR}})$. We note that by Theorem 15 and Theorem 9, it follows that such a system also FTA-complies with the architectural specification $(\mathcal{DG}, \mathcal{C}_{\mathcal{DG}})$.

6.3 Electronic Election

We describe a system that implements the architectural specification $(\mathcal{EE}, \mathcal{C}_{\mathcal{EE}})$ for an electronic election of Section 4.4.

We first select a particular architectural interpretation $\mathcal{I} = (A, \mathtt{dom}, \mathtt{I})$ that satisfies this specification. We suppose that the election is a referendum with the voters voting either "yes" or "no", and the decision determined by a majority of the voters. We take the set of domains to be $V \cup \{\mathit{ElecAuth}\}$, where $V = \{V_1, \ldots, V_n\}$ represents the set of voters, and *ElecAuth* is the election authority. The set of actions $A = \{\mathtt{yes}^v \mid v \in V\} \cup \{\mathtt{no}^v \mid v \in V\} \cup \{\mathtt{tally}\}$, with associated domains given by $\mathtt{dom}(\mathtt{yes}^v) = \mathtt{dom}(\mathtt{no}^v) = v$ for all $v \in V$ and $\mathtt{dom}(\mathtt{tally}) = \mathit{ElecAuth}$. Note that these actions are voter-homogeneous, as required by $\mathcal{C}_{\mathcal{EE}}$.

The architecture has just one edge label *results*. Thus, for the interpretation $\mathtt{I}$, we need to define $\mathtt{I}(\mathit{results})$. Given a sequence α, define the *latest action of voter* v to be the action $a \in \{\mathtt{yes}^v, \mathtt{no}^v\}$ such that $\alpha = \alpha_0 a \alpha_1$ and α_1 contains no action b with $\mathtt{dom}(b) = v$, if such a decomposition exists, or $\perp$ otherwise. We now define $\mathtt{I}(\mathit{results})(\alpha, \mathtt{tally})$ for $\alpha \in A^*$ to be the number of voters v whose latest action is $\mathtt{yes}^v$. (Since the election authority has only the one action $\mathtt{tally}$, this is all that is required to specify $\mathtt{I}$.) That is, in this interpretation, the information that the election authority is permitted to reveal is the number of voters who have voted yes in the latest round of voting. (We assume that a round consists of the events between two consecutive $\mathtt{tally}$ actions, but that if

a voter does not vote in round, their vote defaults to their vote in a previous round, if any.) It is easily seen that this interpretation is identity-oblivious. Thus $\mathcal{I} \in \mathcal{C}_{\mathcal{EE}}$.

Next, we describe a system M with structured state and argue that it FTA-complies with interpreted architecture $(\mathcal{EE}, \mathcal{I})$. We take the set of objects of M to be the set $\texttt{Obj} = \{v_1, \ldots, v_n, bb\}$, where v_i represents the election authority's record of the vote of voter i and b represents a bulletin board where the results of the election are broadcast to the voters. The objects v_i and bb take a value in $\{\perp, Y, N\}$, with $\perp$ indicating that no vote has yet been made by the voter, or in the case of bb, that the election authority has not yet announced a result. In the initial state, all objects take value $\perp$.

Access control on these objects is captured by the following "capability" formatted table:

	V_i	*ElecAuth*
`observe`	bb	$v_1, \ldots, v_n, bb$
`alter`	v_i	bb

It is straightforward to verify that this table satisfies condition AOI′ for policy $\mathcal{EE}$. We let each domain's observation consist of the values of its observable objects, i.e., $\texttt{obs}_u(s) = s \restriction \texttt{observe}(u)$, so that RM1 is satisfied trivially.

The effect of the actions on the state is given by the following code:

$\texttt{yes}^{V_i}$: $v_i := Y$

$\texttt{no}^{V_i}$: $v_i := N$

`tally` : if $|\{i \mid v_i = Y\}| \geq n/2$ then $bb := Y$ else $bb := N$

These definitions satisfy conditions RM2 and RM3. For RM2, note that the actions $\texttt{yes}^{V_i}$ and $\texttt{no}^{V_i}$ change v_i in a way that depends only on the action. The action `tally` changes b in a way that depends on $v_1 \ldots, v_n$, but all of these objects are observable to *ElecAuth*. For RM3, we have that the complement of $\texttt{alter}(V_i)$ is $\texttt{Obj} \setminus \{v_i\}$, but V_i's actions change only v_i, and the complement of $\texttt{alter}(\textit{ElecAuth})$ is $\texttt{Obj} \setminus \{bb\}$, and none of the v_i are changed by *ElecAuth*'s action `tally`.

To show that this system FTA-complies with the interpreted architecture $(\mathcal{EE}, \mathcal{I})$, we use conditions E1-E2 and I1′-I2′. Let the functions F be given by $F(\textit{results})(s, \texttt{tally}) = (|\{i \mid s(v_i) = Y\}| \geq n/2)$, i.e. , the output of F is a boolean value that indicates whether the majority of the v_i have value Y.

We first show that these satisfy E1 and E2. Condition E2 is trivial, since it is never the case that $F(\textit{results})(s, a) = \epsilon$. For E1, we first claim that for all sequences $\alpha \in A^*$, we have that $(s_0 \cdot \alpha)(v_i) = Y$ iff the latest action of V_i in α is `yes`. The proof of this benefits from RM3: since $v_i \in \texttt{alter}(\texttt{dom}(a))$ iff $\texttt{dom}(a) = V_i$, we have by a straightforward induction that $(s_0 \cdot \alpha)(v_i) = (s_0 \cdot \alpha \restriction \{V_i\})(v_i)$. The latest action of domain V_i in α is the final action (if any) in $\alpha \restriction \{V_i\}$, and $(s_0 \cdot \alpha \restriction \{V_i\})(v_i) = Y$ just in case this action exists and equals `yes`. The claim now follows.

We now verify the conditions I1′ and I2′. Condition I1′ is trivial since we never have $F(\mathit{results})(s, \mathtt{tally}) = \epsilon$. For condition I2′, suppose that $\mathtt{dom}(a) \overset{f}{\rightarrowtail} u$ with $f \neq \top$ and $\mathtt{dom}(b) \overset{g}{\rightarrowtail} u$ with $g \neq \top$. Then we must have $f = g = \mathit{results}$ and $\mathtt{dom}(a) = \mathit{ElecAuth}$, hence $a = \mathtt{tally}$ and $u = V_i$ for some i. Suppose additionally that $F(f)(s, a) = F(g)(t, b) \neq \epsilon$ and $x \in \mathtt{observe}(u) \cap (\mathtt{alter}(\mathtt{dom}(a)) \cup \mathtt{alter}(\mathtt{dom}(b)))$ and $s(x) = t(x)$. Then $x \in \mathtt{observe}(V_i) \cap \mathtt{alter}(\mathit{ElecAuth})$, so $x = bb$. By definition of F, we have $F(\mathit{results})(s, \mathtt{tally}) = (|\{i \mid s(v_i) = Y\}| \geq n/2) = (|\{i \mid t(v_i) = Y\}| \geq n/2) = F(\mathit{results})(t, \mathtt{tally})$. In case this boolean value is true, we have $(s \cdot a)(x) = Y = (t \cdot a)(x)$, otherwise $(s \cdot a)(x) = N = (t \cdot a)(x)$. In either case, $(s \cdot a)(x) = (t \cdot a)(x)$, as required.

We remark that in this argument, we have used a function $F(\mathit{results})$ that is not an equivalent state-based encoding of the interpretation $\mathtt{I}(\mathit{results})$, but which is weaker than this interpretation. Correspondingly, in the implementation, the election authority reveals less information to the voters than the architectural interpretation permits. The architectural interpretation permits the election authority to reveal the *number* of voters who have voted "yes", but in the implementation the election authority only reveals whether this number is at least $n/2$. The notion of FTA-compliance with an architecture allows this kind of weakening of information flows in the implementation.

7 Enforcing Architectures on Concrete Platforms

The access control model presented in Section 6 gives an abstract view of how an extended architecture might be enforced. It identifies a set of conditions whose satisfaction suffices to ensure that a concrete system is compliant with the architecture. The process for verification of these conditions in specific settings is likely to be dependent on the particulars of the implementation platform(s) being used. We briefly discuss a few of the possible platforms and the techniques that might be used to show that the access control conditions hold.

Boettcher et al. [2008] survey techniques to achieve *separation* of components, that is, to ensure that communication between components is in accordance with the architecture. In the case of unextended architectures, a very common technique to ensure compliance with the architecture is to map information flow edges to physical causality and use physical separation where there is no edge. Thus, the architecture $\mathcal{HL}$ is commonly enforced in military settings by mapping H and L to distinct processors and/or networks and using trusted devices (data diodes) to ensure a one-way information flow from L to H. The Starlight Interactive Link [Anderson et al., 1996] is a trusted device that can be added to such an implementation to extend it to an implementation of the architecture $\mathcal{SL}$.

One of the longstanding objectives of research on military-grade security has been to avoid the redundancy and consequent expense of such physical implementations, through the use of implementations that enable different security levels to share resources such as memory, processors and networks. An implementation technique that forms the basis for much work in MILS security is the

use of separation kernels, which are highly simplified operating systems with the sole functionality of enforcing an information flow policy.

The key mechanism used to achieve this is typically enforcement of an access control policy by careful management of hardware access control settings and processor modes to ensure that when a process runs, it may read and write only memory regions authorized by the access control policy. Attempts to read or write memory regions that violate the policy are denied by the hardware access control measures. This ensures that the reference monitor conditions RM2 and RM3 are satisfied. RM1 could be guaranteed by ensuring that the hardware access control setting ensures that all peripheral devices with which a user may interact are mapped to memory regions that are associated to the domain of that user. Use of separations kernels introduces the risk that there are covert channels, but much progress has been made in recent years towards formal proofs that separation kernels enforce an information flow policy (e.g., Greve et al. [2003]; Heitmeyer et al. [2006]; Murray et al. [2013]).

An alternative to the use of hardware access control to ensure satisfaction of the reference monitor conditions is to verify, e.g., using static analysis methods, that the code in each domain (e.g., the code describing how the actions of the domain affect the state) reads and writes only locations that are permitted by the policy. Once this has been done, this code can safely run free of hardware access control. This approach is taken in the Singularity system [Hunt and Larus, 2007].

The specifics of the static analysis techniques to be used to enforce the reference monitor conditions will be very language dependent, but can benefit from programming language techniques including type safety and encapsulation constructs including objects, object ownership, and aspects [Kiczales, 1996]. Some of the abstract reasoning in the examples above is already suggestive of such techniques. For example, in the election example, the Election Authority can be viewed as owning the objects v_i and voter V_i can be viewed as having a capability to call a method on object v_i. The code pattern used in the downgrader example for actions $a \in A_T$ could be enforced using aspect-oriented techniques.

Once a basic (unextended) architecture has been shown to be enforced by the implementation, it remains to demonstrate that the trusted components in the architecture satisfy their local constraints (e.g., constraints I1 and I2). As these local constraints are application specific, and implementation dependent, it seems unlikely that a single methodology will suffice. We expect that theorem proving, model checking, and language-based information flow techniques [Sabelfeld and Myers, 2003] may all be used to provide assurance of satisfaction of the local constraints and the filtering requirements introduced by our extended architectures.

We note that our framework is highly expressive, and it may not always be possible to show compliance with a filtering requirement using only local information. For example, directly implementing the filtering requirements of the downgrader architectural specification $(\mathcal{DG}, \mathcal{C}_{\mathcal{DG}})$ by means of controls at the downgrader component D would seem to require the cooperation of High level components (e.g., provision of secure provenance information) to ensure

that the downgrader does not release information concerning C. (The refined architectural specification $(\mathcal{DGR}, \mathcal{C}_{\mathcal{DGR}})$ avoids the need for this cooperation by shifting the trust boundary.) Identifying sufficient conditions for local verification of compliance is an interesting topic for future research.

8 Related Work

The most closely related work is that of van der Meyden [2007, 2012], who defines TA-security and considers refinement of architectures on the basis of this semantics. Our contribution is to show that the definition of TA-security supports derivation of global information security properties, to extend TA-security to architectures that include filter functions, and to develop an account of architectural refinement for these extended architectures. The extension provides a way to specify the behavior of trusted components in a system: intuitively, if a component is the source of an edge in the architecture labeled by a filter function, then the component is trusted to ensure that the interpretation of the filter function limits the information that may flow along the edge. We have presented a number of examples that show that interesting information-theoretic global security properties can be derived from the very abstract statement that a system complies with such an extended architecture and a set of additional local properties. We note that these global properties are more general and application specific than the very particular property "Low does not know any High secrets" that is most often considered in the literature.

Other work has sought to formally describe system architecture (e.g., AADL and Acme [Garlan et al., 2000]), and reason about the properties of systems conforming to a given architecture. There are many software engineering concerns that can be reasoned about in architectural design, such as maintainability, and reliability. This work focuses on reasoning about the information security of systems, and, as such, our architectures specify local constraints on information that may be communicated between components. The local constraints allow the inference of global information security properties. This work is complementary to work on other aspects of system design.

Relatively little theoretical work takes an architectural perspective on information security. One interesting line of work [Hansson et al., 2008] that takes a similar perspective to ours is conducted in the context the architectural modeling framework AADL. This work is based on the Bell La Padula model [Bell and La Padula, 1976]. In a similar spirit are works on Model Driven Security, which extend UML with security modeling notations. Basin et al. [2006] focus on a UML extension for role-based access control policies and model transformations to implementation infrastructures such as Enterprise Java Beans or .NET. Jürjens [2005] extends UML with a focus on reasoning about secrecy in distributed applications employing security protocols. None of these approaches use the application-specific abstract non-interference-style semantics that underpins our contribution, nor do they target the type of reasoning envisaged in the MILS community for development of high-assurance systems built on

infrastructure such as separation kernels.

Standard notions of refinement reduce the possible behaviors of a system. However, arbitrary behavioral refinements of a system will, in general, not preserve information security properties [Jacob, 1989]. Previous work has investigated restricting behavioral refinements to preserve information security properties [Graham-Cunning and Sanders, 1991; O'Halloran, 1992; Roscoe, 1995; Bossi et al., 2003; Morgan, 2006]. Architectural refinement as used in this paper allows only refinements that reduce the information communicated between system components. As shown in Section 5, this notion of refinement preserves certain information security properties. Information security can also be preserved under refinement by modifying the refined system [Jacob, 1989; Mantel, 2001]. Other work distinguishes nondeterminism of a system specification from nondeterminism inherent in the system, and allows refinement only of specification nondeterminism [Seehusen and Stolen, 2006; Jürjens, 2005; Bibighaus, 2006]. These works typically consider only the simple policy $L \rightarrowtail H$, rather than the more general intransitive policies considered here.

Most work on architectural refinement is not directly concerned with information security. Zhou and Alves-Foss [2006] propose architecture refinement patterns for Multi-Level Secure systems development, but do not provide formal semantics. A series of papers [Moriconi and Qian, 1994; Moriconi et al., 1995, 1997] express architectural designs as logical theories and refinement as a mapping from an abstract theory to a concrete theory. This approach is used to establish security properties in variants of the X/Open Distributed Transaction Processing architecture, using the Bell La Padula model [Bell and La Padula, 1976], which lacks the kind of information flow semantics that we have studied here. It is not clear if this approach is sufficiently expressive to represent architectural refinement as used in this paper, and reason about the preservation of information security under such refinement.

Preservation of information flow properties under system composition has, like refinement, been considered problematic. In general, the composition of two secure systems is not guaranteed to be secure. The reason is essentially the same as for refinement: composition reduces the set of possible behaviors of a system, enabling an observer to make additional deductions.

A number of approaches have been developed that allow security of a composed system to be derived from security of its components. These include use of a stronger definitions of security such as restrictiveness [McCullough, 1990] or bisimulation-based nondeducibility on compositions [Focardi and Gorrieri, 1994]. McLean [1996] proposes a framework for specifying and reasoning about system composition, and the preservation of possiblistic security properties. In this work, we are not concerned with showing that security properties that hold of components also hold of a composite system. Instead we are concerned with proving global security properties, and identifying local constraints that components must satisfy. The literature on preservation of information flow security under composition has also largely limited itself to the simple policy $L \rightarrowtail H$.

To some extent, process algebraic operations can be viewed as expressing architectural structure. For example, one could take the view that a process

constructed as the parallel composition of two processes P and Q, with actions A of the composition then hidden, corresponds to an architecture in which P and Q are permitted to interact, but which the environment is not permitted to influence through the set of actions A. However, the semantics of these operators usually do not preserve this view: typically this composition is understood in terms of its possible behaviors with respect to the actions that have not been hidden, and the fact that the system has been composed out of two components permitted to interact in a particular way is lost in the meaning of the composition.

Downgrading has historically been one of the motivations for generalizing the notion of noninterference to intransitive policies. Roscoe and Goldsmith [1999] argued against the ipurge-based semantics for noninterference on the grounds that the meaning it gives to the downgrader policy $H \rightarrowtail D \rightarrowtail L$ is too permissive. According to this semantics, any action by the downgrader D "opens the floodgates," in the sense that it allows all information about the High security domain H to flow to the low security domain L. Roscoe and Goldsmith proposed to deal with this issue by making the semantics of intransitive noninterference significantly more restrictive, in effect reverting to the purge-based definition of Goguen and Meseguer [1982, 1984]. Our approach to downgrading, using a filter function, provides an alternative approach that enables explicit specification of the information permitted to be released by the downgrader.

More recent work on downgrading has concentrated on downgrading in the setting of language-based security. Sabelfeld and Sands [2005] briefly survey recent work on downgrading in language-based settings, and propose several dimensions of downgrading, and prudent principles for downgrading. They regard intransitive noninterference as specifying *where* (in the security levels) downgrading may occur. Since we interpret security levels as system components, our architectures specify where in the system downgrading may occur. The filter functions that we propose in this work specify *what* information can be downgraded and *when* this may occur. Thus, our work combines the *what*, *where*, and *when* dimensions of downgrading. Recent work also considers multiple dimensions of downgrading, including Barthe et al. [2008], Banerjee et al. [2008], Mantel and Reinhard [2007] and Askarov and Sabelfeld [2007b].

Recent work [Askarov and Sabelfeld, 2007a; Banerjee et al., 2008] considers "knowledge-based" approaches to downgrading in language-based settings. However, they do not reason about security properties as general and application specific as used in this paper. O'Neill [2006] uses epistemic logic to specify many information security properties, but does not directly consider downgrading.

9 Conclusion

Through the examination of a number of examples, we have shown that strong information security properties can be proven about a system from a high-level architectural description of the system. Any system that complies with the

architecture will satisfy the information security properties that can be proven about the architecture.

We extended the notion of system architecture to allow finer-grain specification of what information may be sent between components. This enables the proof of stronger security properties, while continuing to providing the benefits of using a high-level architectural description. We generalized the notion of architectural refinement [van der Meyden, 2012] for the extended architectures. Certain security properties are preserved by architectural refinement.

The MILS vision is to build high-assurance systems with well-understood security properties by composition of COTS infrastructure and trusted components. This work brings us closer to that goal by demonstrating that it is possible to compositionally derive strong, application-specific, information-flow security properties from high-level system specifications.

References

AADL. Architecture analysis and design language (AADL). SAE Standard AS5506/A, Jan. 2009.

J. Alves-Foss, W. Harrison, P. Oman, and C. Taylor. The MILS architecture for high-assurance embedded systems. *International Journal of Embedded Systems*, 2(3/4):239–47, Feb. 2006.

M. Anderson, C. North, J. Griffin, R. Milner, J. Yesberg, and K. Yiu. Starlight: Interactive link. *Annual Computer Security Applications Conference*, pages 55–63, 1996.

A. Askarov and A. Sabelfeld. Gradual release: Unifying declassification, encryption and key release policies. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 207–221. IEEE Computer Society, 2007a.

A. Askarov and A. Sabelfeld. Localized delimited release: combining the what and where dimensions of information release. In *Proceedings of the 2007 Workshop on Programming Languages and Analysis for Security*, pages 53–60. ACM Press, 2007b.

A. Banerjee, D. A. Naumann, and S. Rosenberg. Expressive declassification policies and modular static enforcement. In *Proceedings of the IEEE Symposium on Security and Privacy*. IEEE Computer Society, May 2008.

G. Barthe, S. Cavadini, and T. Rezk. Tractable enforcement of declassification policies. In *Proceedings of the 21st IEEE Computer Security Foundations Symposium*. IEEE Computer Society, June 2008.

D. Basin, J. Doser, and T. Lodderstedt. Model driven security: From uml models to access control infrastructures. *ACM Transactions on Software Engineering and Methodology*, 15(1):39–91, 2006.

D. Bell and L. La Padula. Secure computer system: unified exposition and multics interpretation. Technical Report ESD-TR-75-306, Mitre Corporation, Bedford, M.A., March 1976.

D. Bibighaus. *Applying the doubly labeled transition system to the refinement paradox*. PhD thesis, Naval Postgraduate School, Monterey, 2006.

C. Boettcher, R. DeLong, J. Rushby, and W. Sifre. The MILS component integration approach to secure information sharing. In *Proceedings of the 27th IEEE/AIAA Digital Avionics Systems Conference*, pages 1.C.2–1–1.C.2–14, Oct. 2008.

A. Bossi, R. Focardi, C. Piazza, and S. Rossi. Refinement operators and information flow security. In *Proceedings of the International Conference on Software Engineering and Formal Methods*, pages 44–53, 2003.

R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning About Knowledge*. MIT-Press, 1995a.

R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about Knowledge*. MIT Press, Cambridge, MA, 1995b.

R. Focardi and R. Gorrieri. A classification of security properties for process algebras. *Journal of Computer Security*, 3(1):5–33, 1994.

D. Garlan, R. T. Monroe, and D. Wile. Acme: Architectural description of component-based systems. In G. T. Leavens and M. Sitaraman, editors, *Foundations of Component-Based Systems*, pages 47–68. Cambridge University Press, 2000.

J. A. Goguen and J. Meseguer. Security policies and security models. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 11–20. IEEE Computer Society, Apr. 1982.

J. A. Goguen and J. Meseguer. Unwinding and inference control. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 75–86. IEEE Computer Society, Apr. 1984.

J. Graham-Cunning and J. Sanders. On the refinement of noninterference. In *Proceedings of the IEEE Computer Security Foundations Workshop*, pages 35–42, 1991.

D. Greve, M. Wilding, and W. M. Vanfleet. A separation kernel formal security policy. In *Proceedings of the Fourth International Workshop on the ACL2 Prover and Its Applications*, July 2003.

J. T. Haigh and W. D. Young. Extending the noninterference version of MLS for SAT. *IEEE Transactions on Software Engineering*, 13(2):141–150, 1987.

J. Hansson, P. H. Feiler, and J. Morley. Building secure systems using model-based engineering and architectural models. *CrossTalk: The Journal of Defense Software Engineering*, 21(9), Sept. 2008.

C. L. Heitmeyer, M. Archer, E. I. Leonard, and J. McLean. Formal specification and verification of data separation in a separation kernel for an embedded system. In *Proceedings of the 13th ACM conference on Computer and communications security*, pages 346–355, New York, NY, USA, 2006. ACM.

T. H. Hinke and M. Schaefer. Secure data management system. Technical Report RADC-TR-75-266, System Development Corporation, Nov. 1975.

G. C. Hunt and J. R. Larus. Singularity: rethinking the software stack. *Operating Systems Review*, 41(2):37–49, 2007.

J. Jacob. On the derivation of secure components. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 242–247, 1989.

J. Jürjens. *Secure Systems Development with UML*. Springer, 2005.

G. Kiczales. Aspect-oriented programming. *ACM Comput. Surv.*, 28(4es):154, 1996.

H. Mantel. Preserving information flow properties under refinement. In *Proceedings of the IEEE Symposium Security and Privacy*, pages 78–91, 2001.

H. Mantel and A. Reinhard. Controlling the what and where of declassification in language-based security. In R. D. Nicola, editor, *Proceedings of the 16th European Symposium on Programming*, volume 4421 of *Lecture Notes in Computer Science*, pages 141–156. Springer, 2007.

D. McCullough. A hookup theorem for multilevel security. *IEEE Transactions on Software Engineering*, 16(6):563–568, 1990.

J. McLean. A general theory of composition for a class of "possibilistic" properties. *IEEE Transactions on Software Engineering*, 22(1):53–67, 1996.

C. Morgan. *The Shadow Knows:* refinement of ignorance in sequential programs. In *Mathematics of Program Construction*, pages 359–378, 2006.

M. Moriconi and X. Qian. Correctness and composition of software architectures. In *Proceedings of the 2nd ACM SIGSOFT Symposium on Foundations of Software Engineering*, pages 164–174, 1994.

M. Moriconi, X. Qian, and R. Riemenschneider. Correct architecture refinement. *IEEE Transactions on Software Engineering*, 21(4):356–372, April 1995.

M. Moriconi, X. Qian, R. A. Riemenschneider, and L. Gong. Secure software architectures. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 84–893, 1997.

T. C. Murray, D. Matichuk, M. Brassil, P. Gammie, T. Bourke, S. Seefried, C. Lewis, X. Gao, and G. Klein. sel4: From general purpose to a proof of information flow enforcement. In *IEEE Symposium on Security and Privacy*, pages 415–429, 2013.

C. O'Halloran. Refinement and confidentiality. In *Proceedings of the Fifth Refinement Workshop*, pages 119–139. British Computer Society, 1992.

K. R. O'Neill. *Security and Anonymity in Interactive Systems*. PhD thesis, Cornell University, Aug. 2006.

A. Roscoe. CSP and determinism in security modelling. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 114–221, 1995.

A. W. Roscoe and M. H. Goldsmith. What is intransitive noninterference? In *Proceedings of the 12th IEEE Computer Security Foundations Workshop*. IEEE Computer Society, 1999.

J. Rushby. Noninterference, transitivity, and channel-control security policies. Technical Report CSL-92-02, SRI International, Dec 1992.

A. Sabelfeld and D. Sands. Dimensions and principles of declassification. In *Proceedings of the 18th IEEE Computer Security Foundations Workshop*, pages 255–269. IEEE Computer Society, June 2005.

A. Sabelfeld and A. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications*, 21(1):5–19, Jan. 2003.

F. Seehusen and K. Stolen. Information flow property preserving transformation of UML interaction diagrams. In *Proceedings of the ACM Symposium on Access Control Models and Technologies*, pages 150–159, 2006.

B. M. Thuraisingham. *Database and Applications Security: Integrating Information Security and Data Management*. CRC Press, 2005.

R. van der Meyden. What, indeed, is intransitive noninterference? In *Proceedings of the 12th European Symposium On Research In Computer Security*, volume 4734 of *Lecture Notes in Computer Science*, pages 235–250. Springer, Sept. 2007.

R. van der Meyden. Architectural refinement and notions of intransitive noninterference. *Formal Aspects of Computing*, 24(4-6):769–792, 2012.

W. Vanfleet, R. Beckworth, B. Calloni, J. Luke, C. Taylor, and G. Uchenick. MILS:architecture for high assurance embedded computing. *Crosstalk: The Journal of Defence Engineering*, pages 12–16, Aug. 2005.

J. Zhou and J. Alves-Foss. Architecture-based refinements for secure computer system design. In *Proceedings of the Policy, Security and Trust*, Nov 2006.